

Home
Essays
H&P
Books
YC
Arc
Bel
Lisp
Spam
Responses
FAQs
FAQs
RAQs
Quotes
RSS
Bio
Twitter
Mastodon

PAUL GRAHAM

How to Do Great Work

July 2023

If you collected lists of techniques for doing great work in a lot of different fields, what would the intersection look like? I decided to find out by making it.

Partly my goal was to create a guide that could be used by someone working in any field. But I was also curious about the shape of the intersection. And one thing this exercise shows is that it does have a definite shape; it's not just a point labelled "work hard."

The following recipe assumes you're very ambitious.

The first step is to decide what to work on. The work you choose needs to have three qualities: it has to be something you have a natural aptitude for, that you have a deep interest in, and that offers scope to do great work.

In practice you don't have to worry much about the third criterion. Ambitious people are if anything already too conservative about it. So all you need to do is find something you have an aptitude for and great interest in. [1]

That sounds straightforward, but it's often quite difficult. When you're young you don't know what you're good at or what different kinds of work are like. Some kinds of work you end up doing may not even exist yet. So while some people know what they want to do at 14, most have to figure it out.

The way to figure out what to work on is by working. If you're not sure what to work on, guess. But pick something and get going. You'll probably guess wrong some of the time, but that's fine. It's good to know about multiple things; some of the biggest discoveries come from noticing connections between different fields.

Develop a habit of working on your own projects. Don't let "work" mean something other people tell you to do. If you do manage to do great work one day, it will probably be on a project of your own. It may be within some bigger project, but you'll be driving your part of it.

What should your projects be? Whatever seems to you excitingly ambitious. As you grow older and your taste in projects evolves, exciting and important will converge. At 7 it may seem excitingly ambitious to build huge things out of Lego, then at 14 to teach yourself calculus, till at 21 you're starting to explore unanswered questions in physics. But always preserve excitingness.

There's a kind of excited curiosity that's both the engine and the rudder of great work. It will not only drive you, but if you let it have its way, will also show you what to work on.

What are you excessively curious about — curious to a degree that would bore most other people? That's what you're looking for.

Once you've found something you're excessively interested in, the next step is to learn enough about it to get you to one of the frontiers of knowledge. Knowledge expands fractally, and from a distance its edges look smooth, but once you learn enough to get close to one, they turn out to be full of gaps.

The next step is to notice them. This takes some skill, because your brain wants to ignore such gaps in order to make a simpler model of the world. Many discoveries have come from asking questions about things that everyone else took for granted. [2]

If the answers seem strange, so much the better. Great work often has a tincture of strangeness. You see this from painting to math. It would be affected to try to manufacture it, but if it appears, embrace it.

Boldly chase outlier ideas, even if other people aren't interested in them — in fact, especially if they aren't. If you're excited about some possibility that everyone else ignores, and you have enough expertise to say precisely what they're all overlooking, that's as good a bet as you'll find. [3]

Four steps: choose a field, learn enough to get to the frontier, notice gaps, explore promising ones. This is how practically everyone who's done great work has done it, from painters to physicists.

Steps two and four will require hard work. It may not be possible to prove that you have to work hard to do great things, but the empirical evidence is on the scale of the evidence for mortality. That's why it's essential to work on something you're deeply interested in. Interest will drive you to work harder than mere diligence ever could.

The three most powerful motives are curiosity, delight, and the desire to do something impressive. Sometimes they converge, and that combination is the most powerful of all.

The big prize is to discover a new fractal bud. You notice a crack in the surface of knowledge, pry it open, and there's a whole world inside.

Let's talk a little more about the complicated business of figuring out what to work on. The main reason it's hard is that you can't tell what most kinds of work are like except by doing them. Which means the four steps overlap: you may have to work at something for years before you know how much you like it or

how good you are at it. And in the meantime you're not doing, and thus not learning about, most other kinds of work. So in the worst case you choose late based on very incomplete information. [4]

The nature of ambition exacerbates this problem. Ambition comes in two forms, one that precedes interest in the subject and one that grows out of it. Most people who do great work have a mix, and the more you have of the former, the harder it will be to decide what to do.

The educational systems in most countries pretend it's easy. They expect you to commit to a field long before you could know what it's really like. And as a result an ambitious person on an optimal trajectory will often read to the system as an instance of breakage.

It would be better if they at least admitted it — if they admitted that the system not only can't do much to help you figure out what to work on, but is designed on the assumption that you'll somehow magically guess as a teenager. They don't tell you, but I will: when it comes to figuring out what to work on, you're on your own. Some people get lucky and do guess correctly, but the rest will find themselves scrambling diagonally across tracks laid down on the assumption that everyone does.

What should you do if you're young and ambitious but don't know what to work on? What you should not do is drift along passively, assuming the problem will solve itself. You need to take action. But there is no systematic procedure you can follow. When you read biographies of people who've done great work, it's remarkable how much luck is involved. They discover what to work on as a result of a chance meeting, or by reading a book they happen to pick up. So you need to make yourself a big target for luck, and the way to do that is to be curious. Try lots of things, meet lots of people, read lots of books, ask lots of questions. [5]

When in doubt, optimize for interestingness. Fields change as you learn more about them. What mathematicians do, for example, is very different from what you do in high school math classes. So you need to give different types of work a chance to show you what they're like. But a field should become *increasingly* interesting as you learn more about it. If it doesn't, it's probably not for you.

Don't worry if you find you're interested in different things than other people. The stranger your tastes in interestingness, the better. Strange tastes are often strong ones, and a strong taste for work means you'll be productive. And you're more likely to find new things if you're looking where few have looked before.

One sign that you're suited for some kind of work is when you like even the parts that other people find tedious or frightening.

But fields aren't people; you don't owe them any loyalty. If in the course of working on one thing you discover another that's more exciting, don't be afraid to switch.

If you're making something for people, make sure it's something

they actually want. The best way to do this is to make something you yourself want. Write the story you want to read; build the tool you want to use. Since your friends probably have similar interests, this will also get you your initial audience.

This *should* follow from the excitingness rule. Obviously the most exciting story to write will be the one you want to read. The reason I mention this case explicitly is that so many people get it wrong. Instead of making what they want, they try to make what some imaginary, more sophisticated audience wants. And once you go down that route, you're lost. [6]

There are a lot of forces that will lead you astray when you're trying to figure out what to work on. Pretentiousness, fashion, fear, money, politics, other people's wishes, eminent frauds. But if you stick to what you find genuinely interesting, you'll be proof against all of them. If you're interested, you're not astray.

Following your interests may sound like a rather passive strategy, but in practice it usually means following them past all sorts of obstacles. You usually have to risk rejection and failure. So it does take a good deal of boldness.

But while you need boldness, you don't usually need much planning. In most cases the recipe for doing great work is simply: work hard on excitingly ambitious projects, and something good will come of it. Instead of making a plan and then executing it, you just try to preserve certain invariants.

The trouble with planning is that it only works for achievements you can describe in advance. You can win a gold medal or get rich by deciding to as a child and then tenaciously pursuing that goal, but you can't discover natural selection that way.

I think for most people who want to do great work, the right strategy is not to plan too much. At each stage do whatever seems most interesting and gives you the best options for the future. I call this approach "staying upwind." This is how most people who've done great work seem to have done it.

Even when you've found something exciting to work on, working on it is not always straightforward. There will be times when some new idea makes you leap out of bed in the morning and get straight to work. But there will also be plenty of times when things aren't like that.

You don't just put out your sail and get blown forward by inspiration. There are headwinds and currents and hidden shoals. So there's a technique to working, just as there is to sailing.

For example, while you must work hard, it's possible to work too

hard, and if you do that you'll find you get diminishing returns: fatigue will make you stupid, and eventually even damage your health. The point at which work yields diminishing returns depends on the type. Some of the hardest types you might only be able to do for four or five hours a day.

Ideally those hours will be contiguous. To the extent you can, try to arrange your life so you have big blocks of time to work in. You'll shy away from hard tasks if you know you might be interrupted.

It will probably be harder to start working than to keep working. You'll often have to trick yourself to get over that initial threshold. Don't worry about this; it's the nature of work, not a flaw in your character. Work has a sort of activation energy, both per day and per project. And since this threshold is fake in the sense that it's higher than the energy required to keep going, it's ok to tell yourself a lie of corresponding magnitude to get over it.

It's usually a mistake to lie to yourself if you want to do great work, but this is one of the rare cases where it isn't. When I'm reluctant to start work in the morning, I often trick myself by saying "I'll just read over what I've got so far." Five minutes later I've found something that seems mistaken or incomplete, and I'm off.

Similar techniques work for starting new projects. It's ok to lie to yourself about how much work a project will entail, for example. Lots of great things began with someone saying "How hard could it be?"

This is one case where the young have an advantage. They're more optimistic, and even though one of the sources of their optimism is ignorance, in this case ignorance can sometimes beat knowledge.

Try to finish what you start, though, even if it turns out to be more work than you expected. Finishing things is not just an exercise in tidiness or self-discipline. In many projects a lot of the best work happens in what was meant to be the final stage.

Another permissible lie is to exaggerate the importance of what you're working on, at least in your own mind. If that helps you discover something new, it may turn out not to have been a lie after all. [7]

Since there are two senses of starting work — per day and per project — there are also two forms of procrastination. Per-project procrastination is far the more dangerous. You put off starting that ambitious project from year to year because the time isn't quite right. When you're procrastinating in units of years, you can get a lot not done. [8]

One reason per-project procrastination is so dangerous is that it usually camouflages itself as work. You're not just sitting around doing nothing; you're working industriously on something else.

So per-project procrastination doesn't set off the alarms that per-day procrastination does. You're too busy to notice it.

The way to beat it is to stop occasionally and ask yourself: Am I working on what I most want to work on? When you're young it's ok if the answer is sometimes no, but this gets increasingly dangerous as you get older. [9]

Great work usually entails spending what would seem to most people an unreasonable amount of time on a problem. You can't think of this time as a cost, or it will seem too high. You have to find the work sufficiently engaging as it's happening.

There may be some jobs where you have to work diligently for years at things you hate before you get to the good part, but this is not how great work happens. Great work happens by focusing consistently on something you're genuinely interested in. When you pause to take stock, you're surprised how far you've come.

The reason we're surprised is that we underestimate the cumulative effect of work. Writing a page a day doesn't sound like much, but if you do it every day you'll write a book a year. That's the key: consistency. People who do great things don't get a lot done every day. They get something done, rather than nothing.

If you do work that compounds, you'll get exponential growth. Most people who do this do it unconsciously, but it's worth stopping to think about. Learning, for example, is an instance of this phenomenon: the more you learn about something, the easier it is to learn more. Growing an audience is another: the more fans you have, the more new fans they'll bring you.

The trouble with exponential growth is that the curve feels flat in the beginning. It isn't; it's still a wonderful exponential curve. But we can't grasp that intuitively, so we underrate exponential growth in its early stages.

Something that grows exponentially can become so valuable that it's worth making an extraordinary effort to get it started. But since we underrate exponential growth early on, this too is mostly done unconsciously: people push through the initial, unrewarding phase of learning something new because they know from experience that learning new things always takes an initial push, or they grow their audience one fan at a time because they have nothing better to do. If people consciously realized they could invest in exponential growth, many more would do it.

Work doesn't just happen when you're trying to. There's a kind of undirected thinking you do when walking or taking a shower or lying in bed that can be very powerful. By letting your mind

wander a little, you'll often solve problems you were unable to solve by frontal attack.

You have to be working hard in the normal way to benefit from this phenomenon, though. You can't just walk around daydreaming. The daydreaming has to be interleaved with deliberate work that feeds it questions. [10]

Everyone knows to avoid distractions at work, but it's also important to avoid them in the other half of the cycle. When you let your mind wander, it wanders to whatever you care about most at that moment. So avoid the kind of distraction that pushes your work out of the top spot, or you'll waste this valuable type of thinking on the distraction instead. (Exception: Don't avoid love.)

Consciously cultivate your taste in the work done in your field. Until you know which is the best and what makes it so, you don't know what you're aiming for.

And that *is* what you're aiming for, because if you don't try to be the best, you won't even be good. This observation has been made by so many people in so many different fields that it might be worth thinking about why it's true. It could be because ambition is a phenomenon where almost all the error is in one direction — where almost all the shells that miss the target miss by falling short. Or it could be because ambition to be the best is a qualitatively different thing from ambition to be good. Or maybe being good is simply too vague a standard. Probably all three are true. [11]

Fortunately there's a kind of economy of scale here. Though it might seem like you'd be taking on a heavy burden by trying to be the best, in practice you often end up net ahead. It's exciting, and also strangely liberating. It simplifies things. In some ways it's easier to try to be the best than to try merely to be good.

One way to aim high is to try to make something that people will care about in a hundred years. Not because their opinions matter more than your contemporaries', but because something that still seems good in a hundred years is more likely to be genuinely good.

Don't try to work in a distinctive style. Just try to do the best job you can; you won't be able to help doing it in a distinctive way.

Style is doing things in a distinctive way without trying to. Trying to is affectation.

Affectation is in effect to pretend that someone other than you is doing the work. You adopt an impressive but fake persona, and while you're pleased with the impressiveness, the fakeness is what shows in the work. [12]

The temptation to be someone else is greatest for the young. They often feel like nobodies. But you never need to worry about that problem, because it's self-solving if you work on sufficiently ambitious projects. If you succeed at an ambitious project, you're not a nobody; you're the person who did it. So just do the work and your identity will take care of itself.

"Avoid affectation" is a useful rule so far as it goes, but how would you express this idea positively? How would you say what to be, instead of what not to be? The best answer is earnest. If you're earnest you avoid not just affectation but a whole set of similar vices.

The core of being earnest is being intellectually honest. We're taught as children to be honest as an unselfish virtue — as a kind of sacrifice. But in fact it's a source of power too. To see new ideas, you need an exceptionally sharp eye for the truth. You're trying to see more truth than others have seen so far. And how can you have a sharp eye for the truth if you're intellectually dishonest?

One way to avoid intellectual dishonesty is to maintain a slight positive pressure in the opposite direction. Be aggressively willing to admit that you're mistaken. Once you've admitted you were mistaken about something, you're free. Till then you have to carry it. [13]

Another more subtle component of earnestness is informality. Informality is much more important than its grammatically negative name implies. It's not merely the absence of something. It means focusing on what matters instead of what doesn't.

What formality and affectation have in common is that as well as doing the work, you're trying to seem a certain way as you're doing it. But any energy that goes into how you seem comes out of being good. That's one reason nerds have an advantage in doing great work: they expend little effort on seeming anything. In fact that's basically the definition of a nerd.

Nerds have a kind of innocent boldness that's exactly what you need in doing great work. It's not learned; it's preserved from childhood. So hold onto it. Be the one who puts things out there rather than the one who sits back and offers sophisticated-sounding criticisms of them. "It's easy to criticize" is true in the most literal sense, and the route to great work is never easy.

There may be some jobs where it's an advantage to be cynical and pessimistic, but if you want to do great work it's an advantage to be optimistic, even though that means you'll risk looking like a fool sometimes. There's an old tradition of doing the opposite. The Old Testament says it's better to keep quiet lest you look like a fool. But that's advice for *seeming* smart. If you actually want to discover new things, it's better to take the risk of telling people your ideas.

Some people are naturally earnest, and with others it takes a conscious effort. Either kind of earnestness will suffice. But I doubt it would be possible to do great work without being earnest. It's so hard to do even if you are. You don't have enough margin for error to accommodate the distortions introduced by being affected, intellectually dishonest, orthodox, fashionable, or cool. [14]

Great work is consistent not only with who did it, but with itself. It's usually all of a piece. So if you face a decision in the middle of working on something, ask which choice is more consistent.

You may have to throw things away and redo them. You won't necessarily have to, but you have to be willing to. And that can take some effort; when there's something you need to redo, status quo bias and laziness will combine to keep you in denial about it. To beat this ask: If I'd already made the change, would I want to revert to what I have now?

Have the confidence to cut. Don't keep something that doesn't fit just because you're proud of it, or because it cost you a lot of effort.

Indeed, in some kinds of work it's good to strip whatever you're doing to its essence. The result will be more concentrated; you'll understand it better; and you won't be able to lie to yourself about whether there's anything real there.

Mathematical elegance may sound like a mere metaphor, drawn from the arts. That's what I thought when I first heard the term "elegant" applied to a proof. But now I suspect it's conceptually prior — that the main ingredient in artistic elegance is mathematical elegance. At any rate it's a useful standard well beyond math.

Elegance can be a long-term bet, though. Laborious solutions will often have more prestige in the short term. They cost a lot of effort and they're hard to understand, both of which impress people, at least temporarily.

Whereas some of the very best work will seem like it took comparatively little effort, because it was in a sense already there. It didn't have to be built, just seen. It's a very good sign when it's hard to say whether you're creating something or discovering it.

When you're doing work that could be seen as either creation or discovery, err on the side of discovery. Try thinking of yourself as a mere conduit through which the ideas take their natural shape.

(Strangely enough, one exception is the problem of choosing a problem to work on. This is usually seen as search, but in the best case it's more like creating something. In the best case you create the field in the process of exploring it.)

Similarly, if you're trying to build a powerful tool, make it

gratuitously unrestrictive. A powerful tool almost by definition will be used in ways you didn't expect, so err on the side of eliminating restrictions, even if you don't know what the benefit will be.

Great work will often be tool-like in the sense of being something others build on. So it's a good sign if you're creating ideas that others could use, or exposing questions that others could answer. The best ideas have implications in many different areas.

If you express your ideas in the most general form, they'll be truer than you intended.

True by itself is not enough, of course. Great ideas have to be true and new. And it takes a certain amount of ability to see new ideas even once you've learned enough to get to one of the frontiers of knowledge.

In English we give this ability names like originality, creativity, and imagination. And it seems reasonable to give it a separate name, because it does seem to some extent a separate skill. It's possible to have a great deal of ability in other respects — to have a great deal of what's often called "technical ability" — and yet not have much of this.

I've never liked the term "creative process." It seems misleading. Originality isn't a process, but a habit of mind. Original thinkers throw off new ideas about whatever they focus on, like an angle grinder throwing off sparks. They can't help it.

If the thing they're focused on is something they don't understand very well, these new ideas might not be good. One of the most original thinkers I know decided to focus on dating after he got divorced. He knew roughly as much about dating as the average 15 year old, and the results were spectacularly colorful. But to see originality separated from expertise like that made its nature all the more clear.

I don't know if it's possible to cultivate originality, but there are definitely ways to make the most of however much you have. For example, you're much more likely to have original ideas when you're working on something. Original ideas don't come from trying to have original ideas. They come from trying to build or understand something slightly too difficult. [15]

Talking or writing about the things you're interested in is a good way to generate new ideas. When you try to put ideas into words, a missing idea creates a sort of vacuum that draws it out of you. Indeed, there's a kind of thinking that can only be done by writing.

Changing your context can help. If you visit a new place, you'll often find you have new ideas there. The journey itself often dislodges them. But you may not have to go far to get this benefit. Sometimes it's enough just to go for a walk. [16]

It also helps to travel in topic space. You'll have more new ideas if you explore lots of different topics, partly because it gives the angle grinder more surface area to work on, and partly because analogies are an especially fruitful source of new ideas.

Don't divide your attention *evenly* between many topics though, or you'll spread yourself too thin. You want to distribute it according to something more like a power law. [17] Be professionally curious about a few topics and idly curious about many more.

Curiosity and originality are closely related. Curiosity feeds originality by giving it new things to work on. But the relationship is closer than that. Curiosity is itself a kind of originality; it's roughly to questions what originality is to answers. And since questions at their best are a big component of answers, curiosity at its best is a creative force.

Having new ideas is a strange game, because it usually consists of seeing things that were right under your nose. Once you've seen a new idea, it tends to seem obvious. Why did no one think of this before?

When an idea seems simultaneously novel and obvious, it's probably a good one.

Seeing something obvious sounds easy. And yet empirically having new ideas is hard. What's the source of this apparent contradiction? It's that seeing the new idea usually requires you to change the way you look at the world. We see the world through models that both help and constrain us. When you fix a broken model, new ideas become obvious. But noticing and fixing a broken model is hard. That's how new ideas can be both obvious and yet hard to discover: they're easy to see after you do something hard.

One way to discover broken models is to be stricter than other people. Broken models of the world leave a trail of clues where they bash against reality. Most people don't want to see these clues. It would be an understatement to say that they're attached to their current model; it's what they think in; so they'll tend to ignore the trail of clues left by its breakage, however conspicuous it may seem in retrospect.

To find new ideas you have to seize on signs of breakage instead of looking away. That's what Einstein did. He was able to see the wild implications of Maxwell's equations not so much because he was looking for new ideas as because he was stricter.

The other thing you need is a willingness to break rules. Paradoxical as it sounds, if you want to fix your model of the world, it helps to be the sort of person who's comfortable breaking rules. From the point of view of the old model, which everyone including you initially shares, the new model usually breaks at least implicit rules.

Few understand the degree of rule-breaking required, because new ideas seem much more conservative once they succeed. They seem perfectly reasonable once you're using the new model of the world they brought with them. But they didn't at the time; it took the greater part of a century for the heliocentric model to be generally accepted, even among astronomers, because it felt so wrong.

Indeed, if you think about it, a good new idea has to seem bad to most people, or someone would have already explored it. So what you're looking for is ideas that seem crazy, but the right kind of crazy. How do you recognize these? You can't with certainty. Often ideas that seem bad are bad. But ideas that are the right kind of crazy tend to be exciting; they're rich in implications; whereas ideas that are merely bad tend to be depressing.

There are two ways to be comfortable breaking rules: to enjoy breaking them, and to be indifferent to them. I call these two cases being aggressively and passively independent-minded.

The aggressively independent-minded are the naughty ones. Rules don't merely fail to stop them; breaking rules gives them additional energy. For this sort of person, delight at the sheer audacity of a project sometimes supplies enough activation energy to get it started.

The other way to break rules is not to care about them, or perhaps even to know they exist. This is why novices and outsiders often make new discoveries; their ignorance of a field's assumptions acts as a source of temporary passive independent-mindedness. Aspies also seem to have a kind of immunity to conventional beliefs. Several I know say that this helps them to have new ideas.

Strictness plus rule-breaking sounds like a strange combination. In popular culture they're opposed. But popular culture has a broken model in this respect. It implicitly assumes that issues are trivial ones, and in trivial matters strictness and rule-breaking are opposed. But in questions that really matter, only rule-breakers can be truly strict.

An overlooked idea often doesn't lose till the semifinals. You do see it, subconsciously, but then another part of your subconscious shoots it down because it would be too weird, too risky, too much work, too controversial. This suggests an exciting possibility: if you could turn off such filters, you could see more new ideas.

One way to do that is to ask what would be good ideas for *someone else* to explore. Then your subconscious won't shoot them down to protect you.

You could also discover overlooked ideas by working in the other direction: by starting from what's obscuring them. Every cherished but mistaken principle is surrounded by a dead zone of valuable ideas that are unexplored because they contradict it.

Religions are collections of cherished but mistaken principles. So anything that can be described either literally or metaphorically as a religion will have valuable unexplored ideas in its shadow. Copernicus and Darwin both made discoveries of this type. [18]

What are people in your field religious about, in the sense of being too attached to some principle that might not be as self-evident as they think? What becomes possible if you discard it?

People show much more originality in solving problems than in deciding which problems to solve. Even the smartest can be surprisingly conservative when deciding what to work on. People who'd never dream of being fashionable in any other way get sucked into working on fashionable problems.

One reason people are more conservative when choosing problems than solutions is that problems are bigger bets. A problem could occupy you for years, while exploring a solution might only take days. But even so I think most people are too conservative. They're not merely responding to risk, but to fashion as well. Unfashionable problems are undervalued.

One of the most interesting kinds of unfashionable problem is the problem that people think has been fully explored, but hasn't. Great work often takes something that already exists and shows its latent potential. Durer and Watt both did this. So if you're interested in a field that others think is tapped out, don't let their skepticism deter you. People are often wrong about this.

Working on an unfashionable problem can be very pleasing. There's no hype or hurry. Opportunists and critics are both occupied elsewhere. The existing work often has an old-school solidity. And there's a satisfying sense of economy in cultivating ideas that would otherwise be wasted.

But the most common type of overlooked problem is not explicitly unfashionable in the sense of being out of fashion. It just doesn't seem to matter as much as it actually does. How do you find these? By being self-indulgent — by letting your curiosity have its way, and tuning out, at least temporarily, the little voice in your head that says you should only be working on "important" problems.

You do need to work on important problems, but almost everyone is too conservative about what counts as one. And if there's an important but overlooked problem in your neighborhood, it's probably already on your subconscious radar screen. So try asking yourself: if you were going to take a break from "serious" work to work on something just because it would be really interesting, what would you do? The answer is probably more important than it seems.

Originality in choosing problems seems to matter even more than originality in solving them. That's what distinguishes the people who discover whole new fields. So what might seem to be merely

the initial step — deciding what to work on — is in a sense the key to the whole game.

Few grasp this. One of the biggest misconceptions about new ideas is about the ratio of question to answer in their composition. People think big ideas are answers, but often the real insight was in the question.

Part of the reason we underrate questions is the way they're used in schools. In schools they tend to exist only briefly before being answered, like unstable particles. But a really good question can be much more than that. A really good question is a partial discovery. How do new species arise? Is the force that makes objects fall to earth the same as the one that keeps planets in their orbits? By even asking such questions you were already in excitingly novel territory.

Unanswered questions can be uncomfortable things to carry around with you. But the more you're carrying, the greater the chance of noticing a solution — or perhaps even more excitingly, noticing that two unanswered questions are the same.

Sometimes you carry a question for a long time. Great work often comes from returning to a question you first noticed years before — in your childhood, even — and couldn't stop thinking about. People talk a lot about the importance of keeping your youthful dreams alive, but it's just as important to keep your youthful questions alive. [19]

This is one of the places where actual expertise differs most from the popular picture of it. In the popular picture, experts are certain. But actually the more puzzled you are, the better, so long as (a) the things you're puzzled about matter, and (b) no one else understands them either.

Think about what's happening at the moment just before a new idea is discovered. Often someone with sufficient expertise is puzzled about something. Which means that originality consists partly of puzzlement — of confusion! You have to be comfortable enough with the world being full of puzzles that you're willing to see them, but not so comfortable that you don't want to solve them. [20]

It's a great thing to be rich in unanswered questions. And this is one of those situations where the rich get richer, because the best way to acquire new questions is to try answering existing ones. Questions don't just lead to answers, but also to more questions.

The best questions grow in the answering. You notice a thread protruding from the current paradigm and try pulling on it, and it just gets longer and longer. So don't require a question to be

obviously big before you try answering it. You can rarely predict that. It's hard enough even to notice the thread, let alone to predict how much will unravel if you pull on it.

It's better to be promiscuously curious — to pull a little bit on a lot of threads, and see what happens. Big things start small. The initial versions of big things were often just experiments, or side projects, or talks, which then grew into something bigger. So start lots of small things.

Being prolific is underrated. The more different things you try, the greater the chance of discovering something new. Understand, though, that trying lots of things will mean trying lots of things that don't work. You can't have a lot of good ideas without also having a lot of bad ones. [21]

Though it sounds more responsible to begin by studying everything that's been done before, you'll learn faster and have more fun by trying stuff. And you'll understand previous work better when you do look at it. So err on the side of starting. Which is easier when starting means starting small; those two ideas fit together like two puzzle pieces.

How do you get from starting small to doing something great? By making successive versions. Great things are almost always made in successive versions. You start with something small and evolve it, and the final version is both cleverer and more ambitious than anything you could have planned.

It's particularly useful to make successive versions when you're making something for people — to get an initial version in front of them quickly, and then evolve it based on their response.

Begin by trying the simplest thing that could possibly work. Surprisingly often, it does. If it doesn't, this will at least get you started.

Don't try to cram too much new stuff into any one version. There are names for doing this with the first version (taking too long to ship) and the second (the second system effect), but these are both merely instances of a more general principle.

An early version of a new project will sometimes be dismissed as a toy. It's a good sign when people do this. That means it has everything a new idea needs except scale, and that tends to follow. [22]

The alternative to starting with something small and evolving it is to plan in advance what you're going to do. And planning does usually seem the more responsible choice. It sounds more organized to say "we're going to do x and then y and then z" than "we're going to try x and see what happens." And it is more *organized*; it just doesn't work as well.

Planning per se isn't good. It's sometimes necessary, but it's a necessary evil — a response to unforgiving conditions. It's something you have to do because you're working with inflexible media, or because you need to coordinate the efforts of a lot of people. If you keep projects small and use flexible media, you don't have to plan as much, and your designs can evolve instead.

Take as much risk as you can afford. In an efficient market, risk is proportionate to reward, so don't look for certainty, but for a bet with high expected value. If you're not failing occasionally, you're probably being too conservative.

Though conservatism is usually associated with the old, it's the young who tend to make this mistake. Inexperience makes them fear risk, but it's when you're young that you can afford the most.

Even a project that fails can be valuable. In the process of working on it, you'll have crossed territory few others have seen, and encountered questions few others have asked. And there's probably no better source of questions than the ones you encounter in trying to do something slightly too hard.

Use the advantages of youth when you have them, and the advantages of age once you have those. The advantages of youth are energy, time, optimism, and freedom. The advantages of age are knowledge, efficiency, money, and power. With effort you can acquire some of the latter when young and keep some of the former when old.

The old also have the advantage of knowing which advantages they have. The young often have them without realizing it. The biggest is probably time. The young have no idea how rich they are in time. The best way to turn this time to advantage is to use it in slightly frivolous ways: to learn about something you don't need to know about, just out of curiosity, or to try building something just because it would be cool, or to become freakishly good at something.

That "slightly" is an important qualification. Spend time lavishly when you're young, but don't simply waste it. There's a big difference between doing something you worry might be a waste of time and doing something you know for sure will be. The former is at least a bet, and possibly a better one than you think. [23]

The most subtle advantage of youth, or more precisely of inexperience, is that you're seeing everything with fresh eyes. When your brain embraces an idea for the first time, sometimes the two don't fit together perfectly. Usually the problem is with your brain, but occasionally it's with the idea. A piece of it sticks out awkwardly and jabs you when you think about it. People who are used to the idea have learned to ignore it, but you have the opportunity not to. [24]

So when you're learning about something for the first time, pay attention to things that seem wrong or missing. You'll be tempted to ignore them, since there's a 99% chance the problem is with you. And you may have to set aside your misgivings temporarily

to keep progressing. But don't forget about them. When you've gotten further into the subject, come back and check if they're still there. If they're still viable in the light of your present knowledge, they probably represent an undiscovered idea.

One of the most valuable kinds of knowledge you get from experience is to know what you *don't* have to worry about. The young know all the things that could matter, but not their relative importance. So they worry equally about everything, when they should worry much more about a few things and hardly at all about the rest.

But what you *don't* know is only half the problem with inexperience. The other half is what you do know that ain't so. You arrive at adulthood with your head full of nonsense — bad habits you've acquired and false things you've been taught — and you won't be able to do great work till you clear away at least the nonsense in the way of whatever type of work you want to do.

Much of the nonsense left in your head is left there by schools. We're so used to schools that we unconsciously treat going to school as identical with learning, but in fact schools have all sorts of strange qualities that warp our ideas about learning and thinking.

For example, schools induce passivity. Since you were a small child, there was an authority at the front of the class telling all of you what you had to learn and then measuring whether you did. But neither classes nor tests are intrinsic to learning; they're just artifacts of the way schools are usually designed.

The sooner you overcome this passivity, the better. If you're still in school, try thinking of your education as your project, and your teachers as working for you rather than vice versa. That may seem a stretch, but it's not merely some weird thought experiment. It's the truth, economically, and in the best case it's the truth intellectually as well. The best teachers don't want to be your bosses. They'd prefer it if you pushed ahead, using them as a source of advice, rather than being pulled by them through the material.

Schools also give you a misleading impression of what work is like. In school they tell you what the problems are, and they're almost always soluble using no more than you've been taught so far. In real life you have to figure out what the problems are, and you often don't know if they're soluble at all.

But perhaps the worst thing schools do to you is train you to win by hacking the test. You can't do great work by doing that. You can't trick God. So stop looking for that kind of shortcut. The way to beat the system is to focus on problems and solutions that others have overlooked, not to skimp on the work itself.

Don't think of yourself as dependent on some gatekeeper giving you a "big break." Even if this were true, the best way to get it would be to focus on doing good work rather than chasing influential people.

And don't take rejection by committees to heart. The qualities that impress admissions officers and prize committees are quite different from those required to do great work. The decisions of selection committees are only meaningful to the extent that they're part of a feedback loop, and very few are.

People new to a field will often copy existing work. There's nothing inherently bad about that. There's no better way to learn how something works than by trying to reproduce it. Nor does copying necessarily make your work unoriginal. Originality is the presence of new ideas, not the absence of old ones.

There's a good way to copy and a bad way. If you're going to copy something, do it openly instead of furtively, or worse still, unconsciously. This is what's meant by the famously misattributed phrase "Great artists steal." The really dangerous kind of copying, the kind that gives copying a bad name, is the kind that's done without realizing it, because you're nothing more than a train running on tracks laid down by someone else. But at the other extreme, copying can be a sign of superiority rather than subordination. [25]

In many fields it's almost inevitable that your early work will be in some sense based on other people's. Projects rarely arise in a vacuum. They're usually a reaction to previous work. When you're first starting out, you don't have any previous work; if you're going to react to something, it has to be someone else's. Once you're established, you can react to your own. But while the former gets called derivative and the latter doesn't, structurally the two cases are more similar than they seem.

Oddly enough, the very novelty of the most novel ideas sometimes makes them seem at first to be more derivative than they are. New discoveries often have to be conceived initially as variations of existing things, *even by their discoverers*, because there isn't yet the conceptual vocabulary to express them.

There are definitely some dangers to copying, though. One is that you'll tend to copy old things — things that were in their day at the frontier of knowledge, but no longer are.

And when you do copy something, don't copy every feature of it. Some will make you ridiculous if you do. Don't copy the manner of an eminent 50 year old professor if you're 18, for example, or the idiom of a Renaissance poem hundreds of years later.

Some of the features of things you admire are flaws they succeeded despite. Indeed, the features that are easiest to imitate are the most likely to be the flaws.

This is particularly true for behavior. Some talented people are jerks, and this sometimes makes it seem to the inexperienced that being a jerk is part of being talented. It isn't; being talented is merely how they get away with it.

One of the most powerful kinds of copying is to copy something from one field into another. History is so full of chance discoveries of this type that it's probably worth giving chance a hand by deliberately learning about other kinds of work. You can take ideas from quite distant fields if you let them be metaphors.

Negative examples can be as inspiring as positive ones. In fact you can sometimes learn more from things done badly than from things done well; sometimes it only becomes clear what's needed when it's missing.

If a lot of the best people in your field are collected in one place, it's usually a good idea to visit for a while. It will increase your ambition, and also, by showing you that these people are human, increase your self-confidence. [26]

If you're earnest you'll probably get a warmer welcome than you might expect. Most people who are very good at something are happy to talk about it with anyone who's genuinely interested. If they're really good at their work, then they probably have a hobbyist's interest in it, and hobbyists always want to talk about their hobbies.

It may take some effort to find the people who are really good, though. Doing great work has such prestige that in some places, particularly universities, there's a polite fiction that everyone is engaged in it. And that is far from true. People within universities can't say so openly, but the quality of the work being done in different departments varies immensely. Some departments have people doing great work; others have in the past; others never have.

Seek out the best colleagues. There are a lot of projects that can't be done alone, and even if you're working on one that can be, it's good to have other people to encourage you and to bounce ideas off.

Colleagues don't just affect your work, though; they also affect you. So work with people you want to become like, because you will.

Quality is more important than quantity in colleagues. It's better to have one or two great ones than a building full of pretty good ones. In fact it's not merely better, but necessary, judging from history: the degree to which great work happens in clusters suggests that one's colleagues often make the difference between doing great work and not.

How do you know when you have sufficiently good colleagues? In my experience, when you do, you know. Which means if you're unsure, you probably don't. But it may be possible to give a more concrete answer than that. Here's an attempt: sufficiently good colleagues offer *surprising* insights. They can see and do things that you can't. So if you have a handful of colleagues good enough to keep you on your toes in this sense, you're probably over the threshold.

Most of us can benefit from collaborating with colleagues, but some projects require people on a larger scale, and starting one of those is not for everyone. If you want to run a project like that, you'll have to become a manager, and managing well takes aptitude and interest like any other kind of work. If you don't have them, there is no middle path: you must either force yourself to learn management as a second language, or avoid such projects. [27]

Husband your morale. It's the basis of everything when you're working on ambitious projects. You have to nurture and protect it like a living organism.

Morale starts with your view of life. You're more likely to do great work if you're an optimist, and more likely to if you think of yourself as lucky than if you think of yourself as a victim.

Indeed, work can to some extent protect you from your problems. If you choose work that's pure, its very difficulties will serve as a refuge from the difficulties of everyday life. If this is escapism, it's a very productive form of it, and one that has been used by some of the greatest minds in history.

Morale compounds via work: high morale helps you do good work, which increases your morale and helps you do even better work. But this cycle also operates in the other direction: if you're not doing good work, that can demoralize you and make it even harder to. Since it matters so much for this cycle to be running in the right direction, it can be a good idea to switch to easier work when you're stuck, just so you start to get something done.

One of the biggest mistakes ambitious people make is to allow setbacks to destroy their morale all at once, like a balloon bursting. You can inoculate yourself against this by explicitly considering setbacks a part of your process. Solving hard problems always involves some backtracking.

Doing great work is a depth-first search whose root node is the desire to. So "If at first you don't succeed, try, try again" isn't quite right. It should be: If at first you don't succeed, either try again, or backtrack and then try again.

"Never give up" is also not quite right. Obviously there are times when it's the right choice to eject. A more precise version would be: Never let setbacks panic you into backtracking more than you need to. Corollary: Never abandon the root node.

It's not necessarily a bad sign if work is a struggle, any more than it's a bad sign to be out of breath while running. It depends how fast you're running. So learn to distinguish good pain from bad. Good pain is a sign of effort; bad pain is a sign of damage.

An audience is a critical component of morale. If you're a scholar, your audience may be your peers; in the arts, it may be an audience in the traditional sense. Either way it doesn't need to be big. The value of an audience doesn't grow anything like linearly with its size. Which is bad news if you're famous, but good news if you're just starting out, because it means a small but dedicated audience can be enough to sustain you. If a handful of people genuinely love what you're doing, that's enough.

To the extent you can, avoid letting intermediaries come between you and your audience. In some types of work this is inevitable, but it's so liberating to escape it that you might be better off switching to an adjacent type if that will let you go direct. [28]

The people you spend time with will also have a big effect on your morale. You'll find there are some who increase your energy and others who decrease it, and the effect someone has is not always what you'd expect. Seek out the people who increase your energy and avoid those who decrease it. Though of course if there's someone you need to take care of, that takes precedence.

Don't marry someone who doesn't understand that you need to work, or sees your work as competition for your attention. If you're ambitious, you need to work; it's almost like a medical condition; so someone who won't let you work either doesn't understand you, or does and doesn't care.

Ultimately morale is physical. You think with your body, so it's important to take care of it. That means exercising regularly, eating and sleeping well, and avoiding the more dangerous kinds of drugs. Running and walking are particularly good forms of exercise because they're good for thinking. [29]

People who do great work are not necessarily happier than everyone else, but they're happier than they'd be if they didn't. In fact, if you're smart and ambitious, it's dangerous *not* to be productive. People who are smart and ambitious but don't achieve much tend to become bitter.

It's ok to want to impress other people, but choose the right people. The opinion of people you respect is signal. Fame, which is the opinion of a much larger group you might or might not respect, just adds noise.

The prestige of a type of work is at best a trailing indicator and sometimes completely mistaken. If you do anything well enough,

you'll make it prestigious. So the question to ask about a type of work is not how much prestige it has, but how well it could be done.

Competition can be an effective motivator, but don't let it choose the problem for you; don't let yourself get drawn into chasing something just because others are. In fact, don't let competitors make you do anything much more specific than work harder.

Curiosity is the best guide. Your curiosity never lies, and it knows more than you do about what's worth paying attention to.

Notice how often that word has come up. If you asked an oracle the secret to doing great work and the oracle replied with a single word, my bet would be on "curiosity."

That doesn't translate directly to advice. It's not enough just to be curious, and you can't command curiosity anyway. But you can nurture it and let it drive you.

Curiosity is the key to all four steps in doing great work: it will choose the field for you, get you to the frontier, cause you to notice the gaps in it, and drive you to explore them. The whole process is a kind of dance with curiosity.

Believe it or not, I tried to make this essay as short as I could. But its length at least means it acts as a filter. If you made it this far, you must be interested in doing great work. And if so you're already further along than you might realize, because the set of people willing to want to is small.

The factors in doing great work are factors in the literal, mathematical sense, and they are: ability, interest, effort, and luck. Luck by definition you can't do anything about, so we can ignore that. And we can assume effort, if you do in fact want to do great work. So the problem boils down to ability and interest. Can you find a kind of work where your ability and interest will combine to yield an explosion of new ideas?

Here there are grounds for optimism. There are so many different ways to do great work, and even more that are still undiscovered. Out of all those different types of work, the one you're most suited for is probably a pretty close match. Probably a comically close match. It's just a question of finding it, and how far into it your ability and interest can take you. And you can only answer that by trying.

Many more people could try to do great work than do. What holds them back is a combination of modesty and fear. It seems presumptuous to try to be Newton or Shakespeare. It also seems hard; surely if you tried something like that, you'd fail. Presumably the calculation is rarely explicit. Few people

consciously decide not to try to do great work. But that's what's going on subconsciously; they shy away from the question.

So I'm going to pull a sneaky trick on you. Do you want to do great work, or not? Now you have to decide consciously. Sorry about that. I wouldn't have done it to a general audience. But we already know you're interested.

Don't worry about being presumptuous. You don't have to tell anyone. And if it's too hard and you fail, so what? Lots of people have worse problems than that. In fact you'll be lucky if it's the worst problem you have.

Yes, you'll have to work hard. But again, lots of people have to work hard. And if you're working on something you find very interesting, which you necessarily will if you're on the right path, the work will probably feel less burdensome than a lot of your peers'.

The discoveries are out there, waiting to be made. Why not by you?

Notes

[1] I don't think you could give a precise definition of what counts as great work. Doing great work means doing something important so well that you expand people's ideas of what's possible. But there's no threshold for importance. It's a matter of degree, and often hard to judge at the time anyway. So I'd rather people focused on developing their interests rather than worrying about whether they're important or not. Just try to do something amazing, and leave it to future generations to say if you succeeded.

[2] A lot of standup comedy is based on noticing anomalies in everyday life. "Did you ever notice...?" New ideas come from doing this about nontrivial things. Which may help explain why people's reaction to a new idea is often the first half of laughing: Ha!

[3] That second qualifier is critical. If you're excited about something most authorities discount, but you can't give a more precise explanation than "they don't get it," then you're starting to drift into the territory of cranks.

[4] Finding something to work on is not simply a matter of finding a match between the current version of you and a list of known problems. You'll often have to coevolve with the problem. That's why it can sometimes be so hard to figure out what to work on. The search space is huge. It's the cartesian product of all possible types of work, both known and yet to be discovered, and all possible future versions of you.

There's no way you could search this whole space, so you have to rely on heuristics to generate promising paths through it and hope the best matches will be clustered. Which they will not always be; different types of work have been collected together as much by accidents of history as by the intrinsic similarities between them.

[5] There are many reasons curious people are more likely to do great work, but one of the more subtle is that, by casting a wide net, they're more likely to find the right thing to work on in the first place.

[6] It can also be dangerous to make things for an audience you feel is less sophisticated than you, if that causes you to talk down to them. You can make a lot of money doing that, if you do it in a sufficiently cynical way, but it's not the route to great work. Not that anyone using this m.o. would care.

[7] This idea I learned from Hardy's *A Mathematician's Apology*, which I recommend to anyone ambitious to do great work, in any field.

[8] Just as we overestimate what we can do in a day and underestimate what we can do over several years, we overestimate the damage done by procrastinating for a day and underestimate the damage done by procrastinating for several years.

[9] You can't usually get paid for doing exactly what you want, especially early on. There are two options: get paid for doing work close to what you want and hope to push it closer, or get paid for doing something else entirely and do your own projects on the side. Both can work, but both have drawbacks: in the first approach your work is compromised by default, and in the second you have to fight to get time to do it.

[10] If you set your life up right, it will deliver the focus-relax cycle automatically. The perfect setup is an office you work in and that you walk to and from.

[11] There may be some very unworldly people who do great work without consciously trying to. If you want to expand this rule to cover that case, it becomes: Don't try to be anything except the best.

[12] This gets more complicated in work like acting, where the goal is to adopt a fake persona. But even here it's possible to be affected. Perhaps the rule in such fields should be to avoid *unintentional* affectation.

[13] It's safe to have beliefs that you treat as unquestionable if and only if they're also unfalsifiable. For example, it's safe to have the principle that everyone should be treated equally under the law, because a sentence with a "should" in it isn't really a statement about the world and is therefore hard to disprove. And if there's no evidence that could disprove one of your principles, there can't be any facts you'd need to ignore in order to preserve it.

[14] Affectation is easier to cure than intellectual dishonesty. Affectation is often a shortcoming of the young that burns off in time, while intellectual dishonesty is more of a character flaw.

[15] Obviously you don't have to be working at the exact moment you have the idea, but you'll probably have been working fairly recently.

[16] Some say psychoactive drugs have a similar effect. I'm skeptical, but also almost totally ignorant of their effects.

[17] For example you might give the n^{th} most important topic $(m-1)/m^n$ of your attention, for some $m > 1$. You couldn't allocate your attention so precisely, of course, but this at least gives an idea of a reasonable distribution.

[18] The principles defining a religion have to be mistaken. Otherwise anyone might adopt them, and there would be nothing to distinguish the adherents of the religion from everyone else.

[19] It might be a good exercise to try writing down a list of questions you wondered about in your youth. You might find you're now in a position to do something about some of them.

[20] The connection between originality and uncertainty causes a strange phenomenon: because the conventional-minded are more certain than the independent-minded, this tends to give them the upper hand in disputes, even though they're generally stupider.

The best lack all conviction, while the worst
Are full of passionate intensity.

[21] Derived from Linus Pauling's "If you want to have good ideas, you must have many ideas."

[22] Attacking a project as a "toy" is similar to attacking a statement as "inappropriate." It means that no more substantial criticism can be made to stick.

[23] One way to tell whether you're wasting time is to ask if you're producing or consuming. Writing computer games is less likely to be a waste of time than playing them, and playing games where you create something is less likely to be a waste of time than playing games where you don't.

[24] Another related advantage is that if you haven't said anything publicly yet, you won't be biased toward evidence that supports your earlier conclusions. With sufficient integrity you could achieve eternal youth in this respect, but few manage to. For most people, having previously published opinions has an effect similar to ideology, just in quantity 1.

[25] In the early 1630s Daniel Mytens made a painting of Henrietta Maria handing a laurel wreath to Charles I. Van Dyck then painted his own version to show how much better he was.

[26] I'm being deliberately vague about what a place is. As of this writing, being in the same physical place has advantages that are hard to duplicate, but that could change.

[27] This is false when the work the other people have to do is very constrained, as with SETI@home or Bitcoin. It may be possible to expand the area in which it's false by defining similarly restricted protocols with more freedom of action in the nodes.

[28] Corollary: Building something that enables people to go around intermediaries and engage directly with their audience is probably a good idea.

[29] It may be helpful always to walk or run the same route, because that frees attention for thinking. It feels that way to me, and there is some historical evidence for it.

Thanks to Trevor Blackwell, Daniel Gackle, Pam Graham, Tom Howard, Patrick Hsu, Steve Huffman, Jessica Livingston, Henry Lloyd-Baker, Bob Metcalfe, Ben Miller, Robert Morris, Michael Nielsen, Courtenay Pipkin, Joris Poort, Mieke Roos, Rajat Suri, Harj Taggar, Garry Tan, and my younger son for suggestions and for reading drafts.
