

Why We Should Encourage Cheating On Exams

✦ http://www.goldsborough.me/education/thoughts/2020/11/21/23-55-01-why_we_should_encourage_cheating_on_exams/

Peter Goldsborough

Wed Nov, 25 15:56

Cheating is a subjective concept. When a student “cheats” on an exam by stealthily pulling out their phone under their desk and looking up some information to solve a problem, we only call that “cheating” if somebody made up a rule beforehand that forbids it. These rules are not physical laws set in stone by nature. They are entirely artificial. If the examiner changes the rules to make looking up information on your phone during the exam perfectly legal, it is suddenly no longer cheating. If we can settle on this, we should immediately ask: why are the current rules in place and are they really the most *optimal* rules?

When we ask whether the rules that define cheating on an exam, a competition or an interview are *optimal*, we need to talk about what the end of that exam is in the first place. In the tech industry, where interviews are our “exams”, we talk about extracting *signal* or getting a certain *read* from interview candidates. If a particular interview technique allows us to extract a lot of signal on a certain skill, like coding or system design, or a certain quality, like ability to teamwork or manage people, that means it tells us with sufficient precision that the candidate meets our bar in that category.

For example, the signal a math teacher would like to extract – or should be extracting – from a math final is whether his pupils have intrinsically understood the concepts presented in class since the last exam and could apply them to some useful real world task (assuming the purpose of education is to be useful in life, which may be overly idealistic on my part). Instead, the typical rules educators define around academic dishonesty steer the signal into a totally different direction, usually towards providing an excellent read on how well students can stuff three months worth of content into their short-term memory through rote memorization.

That said, not all rules are bad. Some are necessary to provide the right signal. For example, the signal the olympic committee wants to extract from a 100m track-and-field competition is to determine who is the fastest runner, by virtue of their athleticism and training. Since doping would augment the ability of an athlete beyond their natural athletic ability and the olympic committee is not looking for signal on which pharma company can produce the best doping drug, doping is not allowed. This rule perfectly supports the signal the committee is looking for.

We see that the inherent purpose of rules in an exam is to steer the exam towards delivering a particular signal. We define the signal we want to get from the test and we define the rules that will best provide it. Given this, let's take a look at two domains of testing that I'm particularly familiar with – tech interviews and school exams – and reason about whether we should call “cheating” a felony in these circumstances, or encourage more of it.

Tech Interviews

I've been conducting technical interviews for software engineering candidates for a while now and it has really highlighted to me the criticality of finding the right signals from an interview.

Interestingly, the tech industry is quite flexible in terms of interview methodologies. Different companies have come up with different interview practices, all aimed at determining which engineers will be successful contributors to the company and worth the financial, time and human (coaching) investment of hiring them.

The most common form of interview for software engineering candidates is a “coding interview”, in which the candidate is asked to write code to solve a programming problem. These interviews are notorious for requiring textbook knowledge that is rarely used in an engineer's day-to-day job. They're a bit like asking a cook during an interview to apply knowledge from organic chemistry to produce a new dish, when they usually just use ready-made ingredients to do their job. Really everybody knows that these coding interviews are inaccurate proxies of whether somebody is a useful engineer, but they're still very common, for a few reasons. For one, there's an assumption that somebody who is generally intelligent and able to solve these riddles will also be smart enough to solve day-to-day engineering challenges. Another reason is that once in a blue moon you really do need to go down to the organic chemistry level to solve a problem and you'd rather have somebody on your team who can do that.

The rules of coding interviews are relatively strict. We allow candidates the equivalent of looking words up in a dictionary, but we generally expect them to solve the problem from the ground up. Googling for a solution to the problem outright during an interview would be considered cheating, as would be asking their roommate for help. What's funny about this is that what we consider “cheating” according to the rules we've set up for our interviews, is actually what engineers do all the time at their work. If a problem pops up at my job that requires some arcane textbook knowledge I can't remember, I look it up online. When I can't solve a difficult problem on my own, I ask my teammates if they have any good ideas or know of some existing approaches. It's called resourcefulness and it's called teamwork. In fact, somebody who doesn't know how to effectively

research problems on the internet or is socially incapable or too arrogant to ask teammates for help, and tries to solve every problem alone from the ground up, is not going to be an effective engineer. In the end, we get paid to solve problems efficiently, not to prove to ourselves that we're smart.

Never memorize something that you can look up. – Albert Einstein

If the rules of the game lead to behavior that we actually want to see being classified as “cheating”, doesn't that mean that the rules are bogus and not helping us get the signal we want? Yes. Yes, indeed. What's good about the tech industry is that these flaws in the standard interview process are generally acknowledged. Because of this, companies are trying other approaches, such as “take home” problems. These are usually done autonomously, i.e. without an interviewer on the other side. The candidate just submits their solution when they're done. They take a couple hours to solve and you can use the internet, your favorite tools, textbooks etc. If the format were altered to include teamwork with an employee or another candidate, you could even allow them to ask another person for help. Take home problems are an active step in the right direction of extracting the optimal signal in the optimal way during tech interviews.

School Exams

While the tech industry acknowledges the flaws of the “standard approach” to interviewing and is actively experimenting with alternative approaches, our education system's approach to examination is pretty much universally identical, universally stagnant and universally unquestioned. Ever since the dawn of mass market education the typical format of an exam has placed a strong emphasis on the idea that one's memory is the only source of knowledge one should ever use or will ever have available to solve a problem. In most exams, all other resources such as textbooks, other human beings or – god forbid – the internet, are strictly forbidden by the rules of the game. Tapping knowledge from anywhere other than one's brain is an act of cheating.

Using the compass we've utilized so far in this article to determine whether rules are bogus, we can ask what sort of signal an exam with these rules lets us extract from pupils. I would argue that for most exams, the performance of students gives you a read on *how well those students have prepared for the exam*. For the content of most exams, like vocabulary in a foreign language class, natural processes in a biology, physics or chemistry class or historical events in a history class, this means rote memorization. The exam gives us signal on how well the students have memorized the content they knew was due for the exam. For other exams, like a lot of math classes, we get great signal on how much the students have practiced the content. Very often, this means practicing processes at a mechanical level, to the point where they can be applied to new instances of a problem highly effectively. For example, exams will test how well students can take derivatives of functions. You

cannot memorize all derivatives of all functions, but you can practice the mechanical process of taking a derivative for hours on end until you can take derivatives of almost any function, which is what the exam will test. Very rarely do exams test for an intrinsic understanding of the content.

I repeat: *most school exams give an indication of how well students have prepared for the exam.* If we were to ponder for just a moment about what an exam should actually test for – probably whether students have internalized the knowledge, comprehended why it is important to know this and most importantly whether they would be ready to utilize this knowledge in a real world setting – we would be likely conclude that the last thing we'd want our exam to test for, is whether the students acquired this knowledge exclusively for the exam. We'd probably want to prevent our exam from being *gamed*. Unfortunately, it appears to me that very few educators realize the irony of their examination practices in this respect.

There is one more signal our education system tries to extract from the way it tests its students, which is how one student compares to another. Our education system is obsessed with measuring people, obsessed with statisticizing human beings and bucketing them into narrow castes called A, B, C, D and F. The fundamental purpose of an education system should be to give you knowledge, grow your understanding of the world and foster your talents. In no way does a GPA score contribute to this goal in any way, shape or form. It's sole purpose is to compare, so that some decision maker can compare you to somebody else, or compare your school district to another, or American kids to Chinese kids.

Now that we've established that the real signal educators get from an exam is how well pupils have prepared for it and how the pupils compare to kids in Finland, the rules in place for most exams make a lot more sense. If we allowed students during an exam to look up stuff on the internet, then how could we really test that they put in any effort to prepare specifically for this exam? If we let students brainstorm a solution to an exam question with their classmates, then how could we put a number on Peter and a number on Paul and compare them to each other? More importantly, how could we do this for all 250 students taking the exam? What an inefficient way to test students that would be. The education factory would grind straight to a halt. It's much, much easier to call this behavior cheating and keep mass producing graduates quantified and measured according to standardized criteria.

So let's assume that what we call cheating is henceforth allowed. Every exam in school will be open book and teamwork is allowed. Does that mean that the most resourceful student or the best communicator will succeed the most at the exam? No, because the person who rote memorizes all the content and mechanical processes required for the exam will still be faster, potentially completing more of the questions. Will that person fare better one month later when that knowledge needs to be utilized in a real world setting? Probably not, but he'll do a lot better at the exam, because he's better

at gaming it, better at preparing for it – better at providing the signal that the exam is built to extract. This leads us to the same conclusion many kids make, albeit with a slightly more complex derivation: *exams are stupid*. When it comes to being the optimal vehicle to educate people, they really are. They don't teach students to be resourceful and learn how and where to look up information effectively to solve a problem. They don't teach students the value of teamwork or the compounding effect multiple minds have when brainstorming a problem. Yet these are precisely the skills the knowledge worker of the 21st century needs.

Exams should perish. We know that people learn by example, not by exam. In place of exams, test students' knowledge in the real world, on real problems. All exams can and should be replaced by project work. Further, instead of teaching students how to game exams, teach them the tools they will need to succeed in the real world. Teach them how to Google something they don't know. Teach them how to post a question in a help forum. Teach them how to program and solve data science problems instead of memorizing the times table. Teach them how to brainstorm together, how to hear out another person's ideas without interrupting them. Teach them how to communicate a question clearly to another person without wasting their time so they can help you swiftly. Teach them how to synthesize information from many sources instead of memorizing information. I'd hire one student who was taught these skills over ten students who aced their exams, because I don't need a coworker who can prepare for an exam, I need a coworker who can be resourceful and collaborative in solving problems. If that's who we call a cheat, then I'm a cheat and I want a team of the worst cheaters out there.

Conclusion

So what do I mean with “We Should Encourage Cheating on Exams”? My point is that if what we call “cheating on an exam” is actually behavior we want to encourage, then logically we should encourage more cheating. But that's not actually the right solution. The right solution is not to cheat more, but to alter the rules of the exam so that this behavior is rewarded rather than reprimanded. The rules in place for exams are never set in stone, they are not axioms. They were put in place by somebody, thus they can be changed by somebody else. Ultimately, we decide the rules and the rules we decide on determine the signal we extract from an exam. Whatever test we are setting up, be it an interview or a school exam, we should start from the ground up by asking what signal we want to extract and then construct the minimal set of rules that will give us this signal in the most optimal way. It's that easy. It's that hard.