Unlearned Lessons: Why They are so Hard to Learn, and What Could Actually Help

Dr. Sandro Gaycken

Cybersecurity is an old problem, and even though many approaches of the last decade had interesting effects, we're still far from solving it. Self-iterative, dark complexity is in the way—an intriguing new plague of our age—and only high talents in the right places with leeway for real-world experiments can rescue us from being outpaced by authoritarian models of innovation. To build that, we will have to break some rules.

Bad old wine in pretentious new bottles

To many newcomers, politicians, investors, think-tankers, and the media, cyber frequently seems to be a dashing new problem, a few years old at best, with high dynamics and many confusing things happening every which way in tech, politics, offense, startups. Many do not know that it actually is as old a problem as digital technology itself, and a static one at that. Initial concerns regarding technology, defense, and politics began in the 70s, increased during the 80s, only to become a rather niche problem in the 90s and 2000s. During these two decades, there were occasional peaks of attention in the information security community, which mainly occupied the minds of nerds, spies, hackers, conspiracy nuts, and a very few companies. In the 2010s, they received some significant attention for almost a decade thanks to Stuxnet, the F-35, and Snowden. This was also when the term "cybersecurity" became more en vogue, much to the amusement of the old information security community who felt sci-fied by it and who frequently still consider anyone using the c-word a noob (google it!). As a side effect, the new term also hid the history of the problem to many. But no: it's not a new problem. The terminology and the attention changed. The problem in all its beauty, with all its structures, causes, effects, and hundreds of publications, had already existed for roughly 50-years.

Dr. Sandro Gaycken is director of the Digital Society Institute Berlin. He published five scientific monographs along with more than 60 other publications. Dr. Gaycken is an Oxford Martin School Fellow, a program committee member of the Harvard-MIT workshop series on cyber norms, a Senior Advisor for the AI Initiative at Harvard Kennedy School, a research director at the Paris Grand Ecole Le CNAM, and an IEEE permanent reviewer. As an advisor to the German government, he helped create Germany's foreign cyber policy strategy and served as an expert witness in NATO cyber counterintelligence cases, and as director in NATO's SPS program for national cyber defense strategies. Dr. Gaycken also founded his own cybersecurity companies. One is SECURE ELEMENTS, now HENSOLDT CYBER GmbH, which applies high assurance computing concepts from DARPA projects to develop unhackable embedded systems for defense purposes. His latest invention is the company Monarch Ltd., a private military intelligence provider with active cyber skills.

That change in the last decade with regards to the public coverage of cybersecurity was dramatic. The persistence of the media attention caused companies and politicians to finally take a closer look and see what the information security community already knew. In the early 2010s, the IT ecosystem was full of vulnerabilities and only gave very little superficial and peripheral attention to security. At the time, there were few effective remedies at hand and growing dependencies, inter- and intra- functionalities exponentially, and increasing with huge potential negative effects of global strategic proportions. It took a few more years before people took this insight seriously, partly because the IT industry feared regulation and negative market impacts, so it stymied conversation around the topic. But since we have all acknowledged that cybersecurity is a real, hard, and a bad problem, much has happened in the past 6 or 7 years. There is now a lot of regulation, constant senior management attention in the corporate and political worlds, large IT companies invest heavily in improving their architectures and removing vulnerabilities. Additionally, IT-security startups draw numerous investments, a flurry of conferences, workshops, and diplomatic efforts erupt annually, where almost every country in the world and every criminal organization want to own and do offensive cyber.

At this point, it would be interesting to take a brief look at COVID-19, which is another big problem. Fortunately, the reality is much less debatable than the need for cybersecurity, although it is also very complex and with a high potential impact on society. The incredibly thrilling element about COVID-19 is how fast we arrived at multiple effective solutions to this problem. From the first lockdown to the first wave of effective vaccinations was one year. Just one year! In contrast, given that cyber is also a high-pressure problem consuming billions in research and development over the past decade, we have not achieved much; nothing in the way of a "vaccine" or any major improvements.

Of course, regulators, the IT industry, and all those IT security vendors will consider this debatable. New laws have effects and more sectors are regulated. More products cover more vectors. So yes, there was a change, and, yes, it did get better. We did get something. A lot, even, measured in pure quantities. There are tons of regulations, many new products, hundreds of new startups, improvements in architectural security in the big IT substrates. But then again, hacking is still possible. It got more expensive, took more time and skill. But you still get access to any system you desire. Believe me, I have the catalog. And the reason for that is that all these flashy improvements have blind spots, functional shortcomings, or other significant downsides, with many entanglements and leading to so many compromises that no real hard security has been built, although it is technically possible. Let us take a more detailed look.

IT regulations are either specific and then outdated when they come into place, or they are unspecific and unintentionally grooming a lemon market where the cheapest product to compliance is the default for most of the regulated (which generates the kindergarten level of security). Many of them are also in conflict with other regulations or economic and political interests, such as data and consumer protection or rolling out new paradigms like smart grids, smart cities, or the Internet of Things. As a result, they are frequently sabotaged and require tedious negotiations during which other political needs and favors are brought into play, watering down potential outcomes. Acknowledgment of a horizontal function and competent senior leadership, a Churchill of cybersecurity, would be required to solve that problem. Still, the regulator's competence is far from sufficient for that, as is the willingness to appoint something like a Churchill in such an embattled field of competing interests.

Almost all critical national infrastructures have some "cybersecurity" now. The bigger, the better, but even the small ones have taken care of some basics in information security. This has increased the overall level of security; not just anyone can enter an infrastructure with pre-configured passwords of the "12345" type anymore. But while teenage kids are now mostly kept at bay, what most infrastructure suppliers are willing to pay and what the market delivers for that kind of money is a different story. The overall level of security is better than a decade ago, but in terms of attack effort, not by much. It is a well-known problem without a good solution at hand, and presumably another decade of painful trial-and-error ahead of it. The list of problems is simply too long: too much old legacy IT, old copy-paste libraries, proprietary protocols, some of which are from companies that no longer exist, high defect rates with old bugs, structural flaws in baseline security features, bad security engineering culture at the suppliers' and operators, tough requirements on uptime while the program crashes much more easily, too many mixed-criticality systems emerging with new networking paradigms, conventional IT-security with its tolerances and technical perspective not being a natural fit, and so on. No wonder there is no one-size-fits-all miracle cybersecurity product on the horizon (although many try to market theirs as one).

So, what about the new products from the hundreds of startups? Can't they offer a remedy, with allegedly so much innovation going on everywhere? Not really. Most of the "new" products don't learn from history and follow old paradigms like "detection" or "incident management," which have nothing but failed us before and will continue to do so. The vast majority of the flashy new technologies, approaches, and companies are in this camp: old (and bad) wine in new (and expensive) bottles. Oh, and not little of that old wine even being snake oil. Snake oil salesmanship is always highly rewarded in cybersecurity.

Makers of products that follow new paradigms and bring real change but have severe difficulties explaining their approaches. It is notoriously difficult to assess the impact new ideas could have on security and the potential conflicts with the function of the technology that it's supposed to protect. Unfortunately, information security resembles natural science in this respect. As the potential interactions are too complex to make accurate predictions, any evaluation requires a high degree of experimentation and trial-and-error, with many critical voices and arguments from science, evaluators, and competitors. As a result, many of these new solutions under scrutiny will turn out to be less effective than assumed, which many of their inventors know upfront. Yet, given that cybersecurity is still a highly competitive market, literally all these new products and startups promise blue skies and green fields. They must if they want to generate any degree of attention in this very noisy marketplace.

As a result, new ideas are not anything investors, corporate businesses, or government clients will jump at with high enthusiasm. In some cases, new ideas may draw enough attention from venture capital investors that they may get an investment. Still, it is much harder to get this kind of money and also much harder to push an entirely new product into the market against less invasive, cheaper, and more well-known products. And, at least in Germany, some venture capital investors tend to have humor failures if their risk is not rewarded, so they sue the founders to get some money back. This is very common here and not a particularly good way to incentivize entrepreneurship. If a new idea is supported by an authority with a sufficient pedigree (say, from MIT), it stands a better chance in the market, but many authorities have already signed up with the big IT companies for more money and are thus required to support whatever paradigm big IT thinks is best for business.

Most of the serious security improvements of the last, more hectic decade of information security can be located within the architectural improvements in the big IT substrates. Here, a lot of progress has actually happened and radiated into the wider hemispheres of the problem. We have a measure of security in this field, which boasts impressive progress - the exploit and implant market in offensive cyber.

The prices of exploits and implants underlie many dynamics such as market demand, export control restrictions or similar political decisions, shortage of real talent, quality, exclusivity, reliability of the vendor, fittingness to particular problems or chains, the openness

of clients on issues, trust in general, etc. Additionally, successful operative hacking does not require zero days unless high-value targets are of interest. However, the dynamics in this exotic part of the overall information security market show something very interesting.

Before the last hectic decade, finding and exploiting an access-all-areas, free-beer, total-party zero-day exploit in Windows took about three weeks and was a difficult sell in a black market (a white or grey market did not exist at that time) as too much of that was lying around in the first place. Today, finding and exploiting an exploit for Windows, macOS, or many Linuxes takes six to nine months, and it is most likely far from a total-party exploit, requiring a chain of up to ten or more other exploits to actually get full access. Plus, you require a sophisticated implant to do your bidding inside the system, with many security-evasive measures built-in and very skilled operators with a big fat handbook of dirty tricks to deploy all of that. This improvement is due to the big IT companies paying attention to the problem for several years now and throwing a lot of money at improving their baseline security in architectures. This was a highly effective endeavor - and it is visible in how high prices for exploits have risen. A significant strategic security effect of this entire dynamic is the phenomenon that the low end of cybercrime and state-run hackers cannot simply use the same old methods or zero-days for their activities to go after high-value targets anymore. Many of them got more creative and still get in through social engineering or other silly stuff, but most of them had to settle for lower-end targets and leave the high-end targets to the few high-end attackers. This is one bit of good news from the last decade of improvements.

But is it sufficient? Not really. Even if security in some of these cores improved significantly, attacks still exist, and attackers still get in. Again: I've got the catalog. Why? For once, the big cores are simply too big (and continue to grow faster than they can be tested) to ever be without significant vulnerabilities. Even though prices have soared in the global exploit market, there was never a shortage of supplies. The fantastic, single-hit, access-all-areas backstage pass exploit is much rarer than ever before, almost a mystic beast, one to have back-in-the-old-days drinks over, but if you know how to build a chain, you have some money, baseline competence, and contacts, you will be fine. But even if you do not have that, there are still options. IT is a system with many different components, which will never be from one of the big core IT companies. SolarWinds is an example; and the system, sadly, can still mostly be subverted through its weakest link. Some core assets may still stay secure with new architectures and high assurance secure elements, but a smart attacker will still have enough leeway and options to do whatever they need to do. This is well-known and – sorry – a decades-old supply chain issue. And once attackers are in, they can still scale their attacks every which way and wreak havoc up to a strategic level. We do not see open havoc all too often because (also but not exclusively due to exploit prices) all attackers prefer tactics and strategies which enable them to remain hidden deep inside the system, but for those who know, the option is clearly there, anytime. While individual IT companies may be fine with

their reputational concerns and able to blame others and fend off regulation (which is their primary interest in security in the end), the total system security is far from effective. Billions spent and no cyber "vaccine."

Now how can these problems be solved? First, we must acknowledge that cybersecurity is a rather classical case of market failure and a lemon market. Politics would have to step in and regulate. But politics also fail. They're too stupid and too scared to admit it. So, what now? It will help to look at the crux of all levels of market failure from innovation to investment to clients and policy failure alike. I'm a philosopher by training, so I can take the liberty to invent difficult words for key points, and the one that seems to be most fitting in this case is self-iterative dark complexity. Deal with it!

The interesting problem of self-iterative dark complexity

What is it, and how does it apply to cybersecurity? Problems which are (1) highly complex, (2) highly interactive, (3) under constant change or intentional alteration and manipulation with sometimes systemwide effects, (4) epistemically distributed into many different scientific verticals, (5) with many competent actors with biases and competing interests, (6) with many more incompetent actors operating on narratives, biases, paradigms and also with competing interests, and finally (7) with many dark and inaccessible corners, where assumptions, lies and mythical beasts can roam freely, some of which affect fundamentals for the derivation of empirical evaluations – such problems contain self-iterative dark complexity. In these cases, the scientific method fails us in a very specific way; it renders it impossible to collectively build validated knowledge. Since the knowledge system is vertically too complex and consists of too many interactive moving parts and since vertical paradigms are partially based on assumptions about dark corners, there is no solid ground from which to exclude or include hypotheses across verticals (and in many cases also inside verticals) in an empirically verifiable way. Collectively built knowledge cannot spill out of verticals to ever create collective horizontal knowledge across verticals which would be required to make decisions with little uncertainty about impacts.

This seems to be a prevailing problem in many fields of great complexity and with humans involved with interest and impact. Other examples of such fields include financial markets or climate change. But let's take a closer look at cybersecurity from this perspective. How do the individual characteristics of self-iterative dark complexity apply, and what does this imply if anyone would seriously like to solve the problem (assuming we even find someone willing and able)?

First, let us look at the implications of complexity. In the true sense of the term, the most complex element is the IT substrate and options to attack and subvert it through hacking. The limit to how someone can attack and subvert a highly complex, semantic IT system, including elements like design, production, peripheries, supply chain, and naive users, is literally the hacker's creativity. While methods can be categorized along a kill chain or in

MITRE's fantastic ATT&CK methodology, technical implementations of such methods can vary greatly. New methods can come up whenever someone versed in exploitation finds the time to sit down and play. And since computers are language-based machines, the number of attack options can theoretically be determined by Turing's Halting problem: it is infinite, as an infinite number of misunderstandings can be generated, many of which may constitute a security issue. In practice, it may be different due to the limits of even the most creative hacker, but at any rate, it is enough to create genuine complexity in offensive action. This already radiates complexity into the entire system, as the problem information security must solve to stop this infinite number of options for attacks. While this may be solvable for certain methods or categories of attacks or (an interesting architectural paradigm for computer security) in finite machines, it will never work entirely in complex Turing machines. In fact, despite computer science being an engineering science, most researchers agree that large parts of it now resemble natural science, where you have to experiment with the machine to find out how it behaves. And the technical complexity, which is growing exponentially and is already far too massive for any given IT expert (ask them while having strong drinks, and they will admit it), is not even the only complexity to consider. The entire market environment, the use cases, the legal and political, the regulatory environment all add complexity, all can add vulnerability, and all of them are already too big to understand by themselves.

Second, the complexity is interactive. Any change in the technical, industrial, regulatory, political, or the use case environment can create rippling effects across the system at large and end up producing new weaknesses in very surprising ways. In the technical plane of the problem, this is a very fundamental issue with many unforeseeable effects. But even outside the technical sphere, a lot of this can happen. A good example is a diffuse regulation, incentivizing companies to buy cheap security products which have more vulnerabilities than the system they are supposed to protect while enjoying high privileges, thus making it easier for attackers to attack a system. Great job, politics.

Third, the self-iterative element comes into play. Information technology is not set in stone. It changes every day in dramatic ways. Each developer globally adds roughly a hundred lines of code per day to the total system, much of that interacting with other technologies. So while certain anchors may exist, such as trusted cores without significant change, the system at large is under constant change. Much of it can potentially impact other parts, which cannot be anticipated as the range of interactions and use cases cannot be anticipated.

Fourth, the complexity is too much for humans to understand. It is epistemically distributed into many different scientific verticals, which individually already consume the full attention and intellectual capability of even the brightest minds. This applies to computer science, where a hardware security expert cannot fully understand or work on operating system security anymore and vice versa. Still, it is even more difficult when a technical expert is supposed to work with a policy expert or an industry expert. In addition, verticals tend to conflate and

deepen their world and compete with other verticals on their importance because this is how funding impacts the environment. This competition frequently separates them from each other. It sometimes incentivizes them to accept and reproduce certain simplifying narratives of the problem at large if that helps get the upper hand in the competition. In addition, verticals in this highly complex field rarely produce hard paradigms. Others outside the vertical could derive solid knowledge about it in a horizontal perspective, covering the range of reliable insights of the verticals and deriving action items.

At this point, a brief remark about the role of narratives is important. Whenever a field is highly complex and needs to be translated into layman's terms for practitioners, decision-makers, the public, clients, grandmas, or children, a common practice is to use narratives. Narratives are a short version of causally linked facts with specific opinions derived from them and then connected into key messages, sometimes involving unicorns or fairies of sorts to spice things up a bit. The upside of such narratives is that people who are not deep in the respective facts and fields can participate and form an opinion. The downside of narratives is that they actually do that - while essential facts are not included in the narrative and while the narratives narrated in the best way usually win most of the attention, rendering laypeople into victims of that specific perspective of the narrator. The more complex the field, the more options exist to craft narratives and derive opinions. This is hugely important for the next two elements.

Fifth and sixth, the competition of verticals does not stop at science but continues in politics and industry. Competent and less competent actors engage in discourse on these playing fields, ideally establishing a neutral and honest scientifically informed political and industrial debate to arrive at effective solutions. But since they talk across verticals and with many vertically less informed stakeholders, narratives are required. And at this point, the process of knowledge generation fails us again. There are simply too many actors with bad competence or bias but good storytelling skills capable and willing to hijack relevant discourses.

A great example is the "Paris Call," a massive political protest stemming from different IT industry parts, under the rationale that "they know what they are talking about." The members of the Paris Call have political claims and suggestions for many different aspects of our problem, told in a grandmaster narrative, very understandable and, on the surface, politically-correctish. Macron has gobbled it up, as have many other politicians and state institutions. But the entire show is nothing but a vast pile of IT industry lobby garbage, biased manipulative interest, with half-baked arguments as support. Yet hardly anyone is challenging it. Vertical experts are spread too thin and organized too poorly, many of them have contracts with some of those involved and are not willing to bite the hand that feeds them, and many decision-makers don't want to speak out on something they don't fully understand, even though they suspect something fishy. In France, fortunately, ANSSI, the

national authority on cybersecurity, has spoken out against it, as have some others, but with comparatively small voices and no master narratives. This happens all the time. Interested actors create gravity around manipulated narratives, which less competent actors cling on to as it appeals to their level of understanding and as they can hide their lack of competence by slipping under the umbrella of a higher authority with undebatable competence.

The dark element comes into play with the seventh point. Now it gets sexy. In political and scientific discourse, any layperson would suspect that facts can solve this problem of lying and cheating. Except for the dark corners of the netscape are areas that play a crucial role in deriving arguments and are not systematically empirically accessible. The darkest corner and the root of almost all the other dark corners, later on, is the offense. The offense in cybersecurity is such an art and so difficult and complex that almost no one can derive systematic horizontal knowledge about it. It is arcane and highly complex on a technical level in analysis, development, deployment, and operation; it is esoteric and complicated on a human level in actual operations, politics, competencies, economics –and none of it is visible as it is almost always either secret or criminal, so in the shadows either way. The visible part of this iceberg—security testing, hacking conferences, and the occasional leak or detected and analyzed attack—offers only a glimpse into what is happening and mostly one in a rear-mirror of something that is already the past again. And those who are in the know here rarely make it to the other side in defense.

For good hackers, offense is far more profitable, and spies and criminals do not change their field of activities very often, plus they are not allowed to talk details about their previous activities. But as all military thinkers and common sense alike tell us: offense is the base problem of defense. And if you do not know how offense works (because you only see the biased, different, and less competent tip of the iceberg), you cannot do defense. For most approaches, neither fittingness, effectiveness, nor efficiency can be predicted with precise or sufficient reliability or coverage. These aspects can be post-dated for events from the past and predicted for some narrowly defined categories of attacks, but since the complexity offers so many different, non-linear futures, that does not guarantee for the future, so there cannot be a solid statement about our security. Furthermore, if no one sees the real iceberg, the great narrators can invent all sorts of assumptions about what that looks like to support their interests and present their narratives as principles and facts.

So, this is the problem of self-iterative dark complexity. In sum, it renders the process of generating knowledge relevant to decision-making incredibly hard. The common processes of collective knowledge generation and decision-making simply do not work here. It is too hard to find actual anchors for real progress, in the middle of a storm of biases and interests and missteps.

How to solve the problem?

So, what can be done to solve this problem?

Continuous, expensive, offense-informed, and rapid trial-and-error (in actual reality and not just some research lab) could help as a process. But in addition, and parallel to other problems with self-iterative dark complexity, no one wants to make a wrong decision or have made a wrong decision. So many decision-makers confronted with self-iterative dark complexity tend to do one of three things. First, they do not make a decision at all and hide behind other issues or try to push the decision to someone else. This is such common practice that hackers have an abbreviation for it. It is the "OGP"-strategy. OGP means "Other Guy's Problem." Second, they hide behind authorities such as the big IT companies pushing the Paris Call. In this case, they can always say, "but they did it too, and they must be right." This, in turn, creates gravity for the associated master narratives as the decision-makers who should be neutral and open for trial-and-error now have a vested interest that the authorities they signed up turn out to be right. The third option is to make their own decision, which happens mainly in the more peripheral areas of the problem such as advanced research or startup investments. In this case, however, due to public pressure, high costs, heavy competition, high visibility, and great rewards for success or severe penalties for failures (try to get an investment for a new startup if one has failed), the choice that has been made must have been the right one. So, although this field would be one with great outlooks for trial-and-error, there is little tolerance for the "and-error" part of the process, with any aspects of that being denied as long and as hard as possible.

Another option would be getting better and more neutral competence into knowledge-generation and organizing the translation of that knowledge into decisions. But this is no real news. We have known and admired this problem for a very long time. Raising competent cyber forces or cyber strategists is the most pressing urge of all governments and industries. Funnily enough, I know numerous instances where the government secretly thought that the private sector must have been ahead of them because they could hire real experts, while the private sector secretly believed that the government, with all its might and insights must have much better knowledge about the problem. So each one was eyeing the other and trying to copy them only to repeat their mistakes.

The most beneficial entity to solve the problem of self-iterative dark complexity would be a competent and politically powerful government agency. That would be a game-changer. But cyber capacity-building in governments just does not work either.

Why? Because there is a severe talent shortage and tough competition with the private sector with outrageous salary gaps. Of course, most governments have fabulous cyber experts working for them, who have insights into offense, know a broad range of products and approaches, and could lead a strategic effort to success. Unfortunately, it is rarely more than

a handful (very literally!). And a handful is not enough. To address the entire complexity of the problem, you must distinguish lies from facts, and functions require large teams and a collaborative effort with honest support from outside experts, which is just impossible to build. Many of these fabulous experts have changed many things for the better. But it is a drop in the ocean compared to the avalanche of biased and noisy pointless things happening. This problem was also getting worse under recent political attention. Responsible politicians acted responsibly and allocated resources to cybersecurity in government agencies. Much of this went into tech, but much also went into personnel. This, however, turned out to be more of a problem than a solution. Given the ridiculously lower salaries (often factor 10), governments can only hire a low range of talents. But in that area—hundreds of them. Now hundreds of semi-competent cyber people act as described above and pledge allegiance to shortened, biased narratives while building their own little niche inside those. This, however, again creates more issues and, in many cases, only more nonsense instead of empowering those few real high talents in government.

So, by and large, this is a massive problem with wide-reaching effects, but fortunately, it is a problem with a simple solution. Governments must get talented people into their services to support the loyal high talents they already have and eventually find some minor meaningful tasks for the cyber riffraff or fire them. And THIS will actually be a major task for all Western governments in the future because our futures will be more technologized and more complex than they are now and sufficiently large groups of high talents are – at least at present – the only thing that will help.

But how do we get high talents into government service? Simply paying them what would be necessary does not seem to be an option due to the rigidity of government employment rules.

Four other models exist. The first is the model authoritarian countries use. Sadly, it works very well and has solved the problem for them already, creating a rapidly accelerating edge for them in deep tech and high-tech national matters and a strategic erosion of associated capabilities and effects for us. You simply force the high talents to work for you. No startup, no high salary industry career. Government service. Period. Nice and happy government service, so you do not defect. But government service. Period. But that's not a model for us, of course. What are our options?

The second model is the Israeli one. You incentivize young people to come and support the government in return for a world-class education and options to found your own startup later and become a gazillionaire. This model seems to work reasonably well but has its downsides as well. A funny one just came up recently. The young soldiers, relieved from army service and under investment pressure in their startup, immediately apply what they learned to a product and sell it on the global market within months. However, the army's innovation curve is just not steep enough to be around the next corner already, so Israel's enemies

simply buy their startup products to reverse engineer them, in turn, finding and denying Israel's operations around the world. This annoying side-effect led the Ministry of Defense to put cyber exploits under conventional arms control, in turn making every sale a painful, months-long process during which most exploits are being patched and turned worthless and (intentionally) ruining this part of Israel's vibrant startup landscape. Many of these talented young Israelis established defense startups but with mediocre success. A few of them are brilliant, but due to the complexity of the issue, defense is simply a different job. I can tell you this from experience: the almost-last person on earth you want to configure your defenses is a hacker (the very last is the CFO). You want a hacker to look at it, definitely, but not to actually code it. Don't do that.

The third model is the revolving door in two variants. Either contractors come in as hires, working for governments on industry salaries, or they come in for a time and go back after a while, supporting the government temporarily and getting some excellent contacts out of it. Again, both options have negative side effects. With contractors being hired into the government, control and oversight mechanisms do not always work; the collateral damage being incidents like Snowden with massive strategic disruptions among allies. The revolving door expert on the other hand will always be on the lookout for the next super-job and will be tainted with pre-fixed loyalty to his next employer (or at least not wanting to bite any hands they may want to work for in some other future).

The last model we know is outsourcing the problem as a tech issue to the private sector. This can either happen as outsourcing to large companies or startups. Large companies, sadly, are mostly nothing but a copy of government agencies. Big, bulky, slow, uninspired, in parts lazy, no place for high talents, cheap, and of course biased towards profitability, not towards security (which are two very different things in a lemon market). Startups would be in the right place to build external expertise and set up a functioning process as they have the high talents. High talents love startups. But here, government procurement is pure cancer. While strategic government levels may decide that startups are important and must be supported, procurement agencies think procurement procedures are important - and take their time unless they are explicitly overruled. And that is a lot of time—to get into a government contract takes 24 to 48 months. Sadly, many newer approaches like the German Cyberagentur fail to bypass that time frame as it is simply a government mentality thing. Even if legal and procedural levers are set "go," the cautious bureaucrat will still go for the maximum of "cover my ass," especially under new circumstances – and rethink, reconsider, be advised, be reconfirmed, etc. For startups, that simply does not work. New companies with high talents and focused on difficult strategic problems easily burn a six-digit amount of Euros per month. Investors only give you up to 3 million Euros at best. Accordingly, startups' probability of compromising into some detour to a lemon market product is very high. Or, in offense, take a little trip to a less ethically troubled country to sell some products for suitcases of cash

while proper countries take their time to consider and reconsider their options. Or those startups simply die. It happens a whole lot, burning time and talent, trust, and incentives.

So, none of these options seems to be mature just yet. All of them fail either because of the salary problem or the procurement problem (totalitarianism not being an option).

Thus, the lesson to be learned is simple, short, and sweet and doesn't require much theory or explanation. Look at how short this paragraph is! One of those two sets of rules has to be broken to pieces and written anew. Better both. Suppose our societies want to be equipped for the vastly complex challenges ahead of us, which applies to any problem with self-iterative dark complexity. In that case, they must have high talents understanding those problems in a neutral and democratic way, and high talents actually building technical or economic, or political solutions to these problems. Which, by the way, will also provide us with an instant edge against our totalitarian competitors, as our systems will be more flexible, heterogeneous, and adaptable than theirs. But getting there is super hard, incredibly disruptive, and needs to be executed real fast. Badly needed are entirely new government payment schemes and procurement rules or simply more radical outsourcing of state functions to the private industry. Paying higher six-digit salaries, buying deep tech within a few months, and letting private startups handle the technicalities of complex economic and political issues. Nothing less and nothing much else. Undebatable, by the way. In addition, even then, trial-and-error will be required. Self-iterative dark complexity can never be conquered or entirely solved. It can only be surfed by high competence on waves of real-world trial-and-error experiments. This is what we have to build. Or dinosaurs will rule the techno age. And die out once again.