Tech Faculty 2019 lecture given by Professor Luciano Floridi, 17 October 2019

'Digital ethics and fintech'

Introduction

Good evening, ladies and gentlemen. I just came from another conference accidentally that is also happening today in London, and it's the Academy of Engineers. And before I thank the organisers for this opportunity, I'd like to share one bit of news from that conference.

When the same system was used with the engineers and the question was, are you concerned about the speed of development of technology, the answer was quite remarkable, yes. And those are real engineers, not philosophers who don't quite know about what's going on. So, I think that we saw trust, we saw some uncertainty and we saw some concern from the other conference.

It's certainly a good time to do some ethics. So, thank you very much for this lovely invitation and challenge. Because the point was to make this presentation today somewhat relevant to your daily business. So, instead of trying to cover almost any issue in the general world to do with ethics, I will concentrate a little bit more about artificial intelligence, the topic that seems to be getting hotter these days.

Of course, data inevitably joins that as a force for confusion, and in the interplay between artificial intelligence, algorithms, software and data information, the stuff in the engine, trying to understand what the digital questions we are handling today are. And why as we saw quite clearly are they slightly confusing?

I'm not sure whether it's good news or bad news, but they are confusing for everybody, including the philosophers. So, at least you know there is a shared patterning.

The title should tell you where the link will be with some daily activities of people in this room. I will have a couple of slides on Fintech as well, but before I start takin advantage of my half an hour, this is the general map of the talk and the map will appear again. So, you know where we are if you get lost at a certain point.

I will tell you just a little bit more about this new habitat in which we are spending an increasingly higher amount of time.

And then a view of AI as it evolves, and why that is a social problem as opposed to being a marriage. More of this, but you can get the idea, instead of us getting biology and engineering married into some kind of artificial intelligence… Animal intelligence plus artificial intelligence coming together, what we really are doing, and you will not get this from the newspaper.

But certainly, it seems to me a divorce between the ability to solve a problem, perform a task, do something in the world successfully and any need to be intelligent in doing so. I'll repeat this which may be a little bit confusing but remember it's a divorce. And why that divorce is generating some problems in the AI sector, of course, with plenty of data related issues out there.

Then, diving in a little bit on the Fintech moment, so hopefully that will speak to some of your interest and some of your daily problems. And then what ethics can do for us. Why having an ethical approach as opposed to just a purely legal one or good governance, what ethics can do for us. And then the conclusion.

New habitat

So, as I say, the map will appear again, but let's talk about the new habitat. This is what we are bombarded with every day. Gartner predicts digital ethics and privacy as a strategic trend for 2019. If even Gartner gets it, then obviously everybody does. You don't have to be a philosopher knocking the door of any company or any government that says, we have a problem.

And then you start again and again and again and finally, today, the door is open, and they say, can you come in, we have a problem. Oh, finally, yes we do have a problem.

There is a lot of money going into solving the problem. That is what we can only dream of on this side of the Atlantic, one billion dollars spent on starting AI ethics. If you scratch a little bit, actually that's not true. But of course, as all things that seem to be a little bit too much- do we really need a billion dollars to do this?

In the history of philosophy, we've never seen that much money. So, my students should not have any delusions about what sort of profitable career they're going to have. The truth is that it's going to be a lot of buildings and maybe a couple of faculty hires. But still it is an issue that is attracting a lot of attention.

Why do we have this problem today? I made a form which I hope will help. So, the digital technologies, science, policies, habits, you name it, anything that has something to do with digital stuff has this amazing ability to cut and paste. So, unglue and glue together something we have inherited from the past that we thought was not doable. Or, so solidly united that you couldn't unglue it.

I'll give you a few examples which I hope will speak to you quite clearly.

But imagine for a moment, allowing while I'm talking, there's a time a couple of centuries, we call it modernity, a couple of centuries ago, long story… And here comes this new technology and takes modernity and starts pulling things apart or putting them together in a way we have never seen before. That is something that is extraordinary.

For example, presence and location. This is something that we take for granted today, even people with some white hair like me. But until this generation, you were present with your body physically somewhere, and located there in terms of what you could interact with in one go. It was glued together completely.

Today, of course you shop online while you are at home. So, it took a philosopher to think about it. It's a huge difference. That's why, for example, branches or brands are closing down. This is just in America, and it's already old data, but they keep closing them. Why? Because they have been built at a time when location and presence were one thing. You had to go to the bank with their particular cheque. No longer. No one needs the bank, the bank closes down.

So, when you see a shop or a bookshop store opening a cafeteria, a coffee shop, they're not trying just to make some money on the side. They're trying to re-glue together the presence and location. They just want to make sure there's someone who is there maybe for a cappuccino and then they can browse the books.

So, is there any need for anyone to go to a bank today? If not, you'd better close down the branch, because it won't be useful.

I'll come back to that in a moment, but staying with ungluing, law and territoriality. For about four centuries, since Westphalia almost, that's 1648, top corner on the left, we decided my place, my rules, your place, your rules. It was as simple as that.

It's basically what happened with the Three Musketeers, do you remember vaguely? When they meet again 20 years later, Westphalia has happened, and there's peace in Europe, more or less. Because my place, my rules. Your place, your rules. Place.

So, law became one thing with territoriality. It stops at the border. Now, cyber space? Try that when discussing with Google. I was on the advisory board for the Right to be Forgotten. And, try to tell them, look, France would like you to remove some links from your Google.com. I don't think so. Oh, we like the Three Musketeers. I don't care for France. That's not the point. We are not at this time anymore, and therefore it doesn't work.

So, fast forward a few centuries, and this ungluing of law and territoriality is what people had in mind when they started buying GDPR.

General Data Protection Regulation, what do they find? Well, guess what? I don't care where you process the data, as long as data that belongs to your citizens. And so, trust me… Two or three governments ago when they came to change quite quickly, we were discussing about Brexit. And in context, it was GDPR won't apply. No, it will apply. It doesn't matter, because it's about European data. It's no longer about the state of reality. It's not my place, my rules, your place, your rules. So, we bypass that undoing with GDPR thanks to the fact that we just didn't care about where it happens but to whom it belongs.

Currency has also had its challenged moments. Libra is an attempt to attach currency for money. What's the difference? Well, money at some point was salt, hence why salary. Your salary is because you were paid in salt. That's also why you take things with a pinch of salt, by the way, and I'm not joking. That comes from the fact that you were paid in salt and so a pinch was extra. It's a long story for another time.

But you can pay money equal to Avios on British Airways which would not accept cash cards. So, how do we detach this? Well, Libra, if you know the story is an attempt to work on this detachment.

I would also link information and identity. For any European legislation, we are data subjects. That is what defines us legally. And that growing of my identity, my personality, who I am, information about me is totally new.

No one speaks any more of going online or being offline. If you do, you are from the 90s. We are all onlife these days, because you coin the word. I have used it for some time, companies have found it very helpful. And today it means that it's a mix and match. Because while we're here, if we have our mobile phones on because we want to view the slide and so on, they are, as we know, geolocating each of us several times per minute per app. And so how much offline are we? Not very much really.

So, asking the question, are you online, are you offline? That's when we used modems. Remember those things? Some of you that don't, because I can see that you're young, but some of us do. The whale singing is something that my students this year, they never heard. They don't know what I'm talking about. Yes, in the 90s we used to have a modem, and the whale singing was handshaking between two computers.

We don't do that anymore, and we live in this mixed space of an infosphere where essentially data, life, my kitchen, talking to gadgets… Now, until recently, you were crazy if you were talking to something. The guy is talking to a gadget.

Of course, now what else? That's how I open the window. So, this is a completely different world. And a metaphor here. This is not Oxford yet, but we're getting there with global warming. But the point is that we're living in this kind of a mangrove society. It's a bit of a forest, it's a bit of a mix and match.

In the mangroves, they grow where the very salty water of the sea and freshwater of the river mix. Asking there whether the water is salty or fresh means not having understood where you are. Because the water there is brackish. I had to check on the Oxford English dictionary. That's not the kind of English you learn at school, but it's called brackish water for a reason. It's mixed and it's only those plants that can survive in that kind of water.

Now, imagine societies living in different corners converging around the mangrove place. That's why it's becoming more and more popular. By now, there are so many people online, we have so much data. We have so much computational power, no wonder this is generating new issues, or revising old issues in new formats.

For example, and speaking of banks, this was old news. I'll show you something more recent in the next slide, but this is 2018, and it's Lloyd saying that it's sending home 6,000 but hiring more than 8,000. And then everyone says, that's great, that's fine, we got 2,000 extra people. But they are different.

The 6,000 plus sent home are probably not overlapping with the 8,000 plus hired. So yes, there is a plus here but there's a generation that's paying the cost of this transformation.

So, sometimes when people say, oh my God, they steal our jobs. Well, first of all, they're not migrants. They're not stealing the jobs, but this AI and all this technology is not destroying jobs. It's reshaping them enormously, drastically at their level. And so that's okay. No, it isn't because someone will be paying the price for this transformation.

It’s for another day, another context but that’s where politics plays a role, because we need to make sure that the enormous benefits coming potentially in the future can be spread a little bit to the present, so that those who are paying the price for this transformation are not paying the full price and receiving no benefit. That’s not the kind of society we want to live in.

So yes, some sort of welfare, some support, some way of bridging the gap will be inevitable if we want not to increase unemployment this time. But this is also the world in which we live, and we say, okay, it’s digital.

And we move from saying something is it’s just a fashion, totally useless, it won’t make any difference, to it’s the end of the world.

Well, look at what has happened with Sainsbury. And I know you follow the news so you know what happens next. This was in 29th April this year, and Sainsbury tried this. Remove all the tills, see what happens. People will love it. No, they don't. Things were a disaster.

Welcome to mix and match, it’s always a bit of both. It’s analogue and digital. I like to have someone checking things for me, and yes, I don't quite know how to scan this and scan that. And if there’s a queue, I will have to go there and say, sorry, how do I do this?

So, we need to be careful by swinging between two extremes. It’s going to be nothing, it’s going to be everything. It’s the boring middle where policy happens, and that is exactly where we need to do a lot of work. Now, this is not exciting. No one will ever have a revolution for the boring middle, but that’s where the difference happens. Daily, systematically, step by step.

AI: a divorce

Now, this is the habitat, the divorce that I told you about. Well, why a divorce? And I know we focus more on the technology. If this is the general background, AI has been described in a thousand different ways. And I’m not going to bother you with the details.

In a different life, for the first half of my life I was a mathematical logician, so I dealt with these kinds of figures from the other side of the production line. But today, what we really have is, to use a very trivial example, something that plays chess well, better than anyone else in this room… My phone down there, much to my annoyance, by the way.

And the intelligence, don't tell me that an iPhone is in any way remotely close to a dog, not only a human but in intelligence, a dog. Fine, it plays better chess than anyone in this room. So what? We found a way to make sure that you can play chess without any intelligence. As simple as that, like my dishwasher. It does the dishes better than me, at zero intelligence. And on and on and on.

The extraordinary thing is that we move from doing mechanical things which we were very good at in the past century onwards, to things that would require intelligence if done by a human. If I went to do that myself, I would have to use my brain. That’s not including the dishes, and I’m in charge of the dishwasher, by the way in the house. I want that on record. I’m in charge of the dishwasher.

And that dishwasher is a quintessential example of a problem we have sometimes with AI. Simple example, you come to my house and you find the dishes clean on the table. The question is who did the dishes? The dishwasher or the person? I don't know. Are they clean? Hopefully.

I say, oh my goodness. You can't recognise who did them, therefore the dishwasher, is it intelligent? Of course not.

Just because you cannot reverse from the output the process to identify the input, so the irreversibility of the process does not guarantee that the same process has been used to get the output. So, the clean dishes have been achieved in two different ways. One is mechanical and involves lots of soap and lots of water. A whole box. The other one is one by one by hand. Process different, same result. It doesn't mean that therefore the sources are identical. That's the normal problem.

So, if this divorce is what we are seeing increasingly happen under our eyes, first there is a bit of advice and context to your average corporate start-up company. Where do I find an AI application? Start thinking about divorce. Is there any process in your company where people are required to be intelligent in order to do it, but it could be done without intelligence? Like what?

For example, when was the last time you had an automatic robot filling the car for you at a gas station? Think about it. We have more than a billion cars in the world. They've all been designed to be filled up by a human brain who picks it up, sticks it in, waits and puts it back. That's the deal. That's extraordinary. The most amazing, complex thing ever created by this universe which is a human brain used as an interface between the pump and the car and back. Bad design. We didn't think about it.

Oh, by the way, we're doing the same mistake often enough when it comes to electric cars. Pick it up and stick it in, really? So, can we have something more intelligent here? I don't want to recharge my car. I don't want to be the interface between the two.

So, the divorce for example there would be this can be done without any intelligence. It can be done by a robot. So, there are a billion cars to be served. How many petrol stations would like to buy your robot arm that does that for you? Just in case… 1% of the profits of the company if you open it please, thank you.

AI therefore generates extra challenges because of this divorce. Like what? Well, here's a simple graph and hopefully it will be helpful to carry on the conversation. It could be used where? There are plenty more opportunities. I'm going to list them quite quickly, but it can enable self-realisation. It can enhance our human agency by supporting us. It can ease our societal or group capabilities and it could calculate social cohesion. It could basically help us to be better human beings in a better society.

The same however, turning to red, of all the uses for AI when you devalue human skills, when you remove human responsibility, when you reduce human control or erode human self-examination. It's the same point, but the other side.

And then there is an amber corridor, which normally people seem to be underestimating. And I'm quite keen on stressing that, because the green and red are more or less clear but people forget

about the amber. All the things we're not doing because we don't have a framework that helps us to understand what the right thing to do is.

And by the way, remember one of the answers? Answer, lack of clarity about what the right thing to do is. That's exactly the amber point, I don't know but then we say, sorry, I'm not going to do it. I'm going to look for someone else to do it. Meanwhile, say the NHS, we're not decreasing pain or suffering or even saving lives because AI is something that we're not quite sure of.

So, better safe than sorry does not allow you to leave the cave or use fire in the first place. So, we'd better be careful about using that.

How are we going to handle this?

AI ethics

Well, in the past only the bottom three, the smaller three Prussian dolls were in debate. You had customers, you were constrained by business, which was constrained by law and those were the three interplaying roles. Today, more and more you hear people talking about ethics. It's more a sense of what is socially acceptable.

How customers feed back into the loop and say, look, I'm not quite sure we like this. So, imagine law constraining what business can do which basically tells people what the options are, and then people go back and say, can we just change the law because we don't like it? Or can we have that as a law instead of the other? Can we? So, the socially preferable to the socially acceptable, the red line, gets back into the cycle. Our society can have a big input on the whole process.

Another point, and I know that I'm talking to experts here, there is a bit of interplay between ethics, regulations and governments. So, ethics is very much about what should be happening, what should be done. Regulations is about what may be done and that's different. That's what is allowed or is not forbidden as the best legislation, as opposed to what ought to happen.

So, in terms of regulations on the road, of course that speed limit can be anything else, but if you see someone, maybe an accident, maybe it's a good idea to stop and check. Maybe that person is injured. It's a moral thing. There is no regulation about the road that says stop every time you see someone in a difficult circumstance and so on.

And what about governance? Well, governance can also be more than just ethics and regulation.

It can be the kind of rules the company has internally. Our own regulation is not a legal requirement and it's not about what should or should not happen ethically speaking. Now, mind that all these three do interplay with each other in quite a confusing way.

Ethics, if you like, is about what is acceptable or preferable in society. What's right. Regulations constrains behaviour in terms of compliance. And compliance is absolutely necessary in places where the rules are good. Not in places where the rules are not good. If you are a South African in apartheid time, compliance is not the right thing to do. So, ethics overrides regulations.

That's normally what my law faculty colleagues don't quite like, but it's the case. It's because of ethical reasons that we change the law, not the other way around.

In terms of compliance, I say necessary. I didn't say sufficient. The analogy here is with the rules of a game, say football. Do you want to play football? Those are the rules. Okay that's it. No, maybe I would like to win so I'll just play by the rules. That's not enough. That's insufficient. Otherwise I would be the best tennis player in the world. I obey all the rules. Or with chess or anything else.

So, the other half, more than above compliance is what accepts the necessity of the rules and tells you that more needs to be done. More in what sense? Well, here we are getting close to the ethical discussion. And for non-English speakers sometimes… I find this quite often in Brussels, because there is an S at the end, you hear people talking about ethics are… Like nails down a blackboard, when there were blackboards in the past. It's terrible, ethics are… No, ethics aren't because ethics is.

It's like mathematics, mathematics is, it just happens. For foreigners like me, always a headache. But it comes with an S and it's not plural. That's an exception.

But I like to remind people that no, we will never agree ethics are… No, they are not. Ethics is… We do agree on a number of things, and we have been increasingly agreeing on what's right and wrong. Not that we do the right thing. No, we keep doing the wrong thing, but at least we know what's right and wrong.

Like what? Like this. So, this is a quick list of the most recent influential code of ethics for AI around the world. And it starts from AI UK, that is the House of Lords and so in the past, all the way down to Beijing. And the last three, the European Commission, high level expert group on ethics of AI… I'm a member. Or the AI for People, which is also a project that I chaired for Brussels, for the European Commission. Also, the OECD, the Beijing AI principles.

So, there is some convergence. It doesn't mean that we do the right thing. It doesn't mean that everybody abides.

But it does mean that at least we have some kind of common language to understand what's right and wrong.

And so, beneficence, non-maleficence, autonomy - ours not theirs, not the machines. Our own autonomy preserved, making sure that it's not undermined. Justice or fairness, as you hear often. That's where the unfair bias is on place alone and explicability, understanding what's going on there. I don't want to talk to a black box and say, oh sorry it's just magic. If it's magic, don't use it. Either you know what you're talking about or you don't deploy it.

So, those are points of convergence that should give us some hope for a better future, where essentially more or less we start getting close to something we can all agree about. Once again, not delusional about how well behaved we shall be. It's a matter of understanding.

And how do we translate this into practices? Well, this is where I'm getting close to almost the final part of my talk. We don't have much time in many contexts. Now, technology can help us a lot for example with one of the most pressing issues of our generation, which is global warming, right?

There won't be a future if we don't do something about it. AI can be a great force for good. Remember the divorce? It's a lot of the muscles out there that we can deploy to do the right thing.

We also use it a lot in complex context like big cities, and just statistics about where we are going to be, more and more people in London and all over the place, with tens of millions of people living in the same place. Without AI we won't be able to manage this successfully.

Fintech

But also Fintech, and let's now dive in and I'll come back to the conclusion. Fintech is increasingly automatic. It's increasingly relying on software to do things better, more cheaply, more efficiently, to be more efficacious and to also employ less people. So, it has essentially two functions. It transforms money into a commodity, and it is transforms risk into cost. I know that you do this on a daily basis, so I'm not going to tell you more about this.

But what happens when we transform all this into a digital environment, with plenty of digital data, plenty of digital software and digital resources and AI?

Well, all this… The money to commodity, the risk into cost starts getting to increasing values for everybody, not just the customers, not just the people who offer this. More choice, less about transforming money into a commodity and offers. It means in the Western eyes that you have more freedom. And more certainty. That's what is transforming risk into a cost. You can afford more, you can take more risk, there's more certainty out there. So, more security. All this based on trust, and we saw this and how important for people in this room it's always been.

So, at this point putting all this together, remember all those long lists of principles all over the world? Well those, the A group as it were, can be transformed into something that guides Fintech. So that the green and the red and the amber have a different role in transforming the values of the industry.

Essentially you go from A to B so that C is only green when it has an impact on D. I didn't want to show you this at the beginning because that would have been scary, but I hope this is quite clear at this stage. To do a good job, we need those principles and we need to apply them, we need to make sure that the green takes off and the red is minimised, and the amber is avoided.

Soft ethics

Now at that point, this is policy and what's the conclusion? Policy, forgive me but I've become increasingly fond of it, and not just because we're going to have a policy expert next, but it's what makes the difference. It's design of the rules.

So again, this is what the market value of the top 500 companies looks like, and it's the most recent number we have. It's all about intangible. Intangible means a lot of ethics, goodwill, your name, your reputation, what people will think of your next product.

Clearly, there is a lot of ethics there. What kind of ethics? Ethics that support compliance because we are trying to find a place where the rules are okay, otherwise I would not recommend it.

We're not some cowboy country, we are the UK, we are in Europe. The rules are fine, we can do compliance and be proud of it. Human rights of course. But then feasibility is open, and we have grey areas, empty areas and fillable areas. Remember? We all want to play to win.

Playing according to rules is necessary and probably insufficient. That's where ethics come in and in those principles, in those applications we can make a difference. Signalling, putting in your company that you are serious about the ethics of this. Collaborating with other companies on the same platform. Showing awareness of issues that you have. And maybe also solving some problems.

All this has a twofold value or force. It can be at the same time a risk management issue, if we have the right ethics in place. You can avoid some risks and you can minimise others. And avoid the amber, the opportunity cost.

But it's also an opportunity strategy. The opportunity cost there will not be incurred because of the value you can have.

And finally, so what we need to do here is four points. We have been working on feasible technologies all our lives. We think that something has to be sustainable. Well, certainly given the world state, the least that any company should do is sustainability. Acceptable and preferable socially speaking, when I comes to digital technology, that is increasingly expected from everybody.

So, what we should do and it's the last slide, we should think more deeply about all this and, forgive me as it's bad English, but it sounds better on the slide, more deeply, not deep. We should think more deeply about this. Not do and then think afterwards, no. How many steps do we need to take back in order to jump forward?

We should be more mindful, not careful… Mindful. Mindful of the extra consequences of what happens if…?

We should care more about everything, not human-centric, but world-centric, because we cannot afford to be human-centric anymore. We need to be centric about the whole system here.

And design better, for humanity and the environment. Thank you so much.