We need a personal digital advocate

Mark Braverman[1]

Abstract

There is a growing power imbalance between companies who have access to powerful algorithms and processing capacity and users who don't. To restore the balance, we need to put equally powerful algorithms in the hands of individuals. There are two concrete steps we recommend:

(1) there should be a new regulatory framework which creates an unbreakable commitment for an advocate (a digital service) to work exclusively in the interest of its client;

(2) existing and new regulations around digital rights of individuals (such as GDPR) should make it a priority to make it easy for users to take advantage of these rights using software.

[1] Department of Computer Science, Princeton University. Opinions in this piece are the author's, and do not reflect the views of Princeton University or of funding agencies.

The author thanks Anna Braverman for the many discussions and for significant contributions to the essay's writing.

The luxury of leaving no footprints is getting more and more elusive as the places we walk get more and more virtual. We can by all means be incognito when we visit a physical space (though even this is getting eroded as commercial surveillance gets cheaper are more ubiquitous), but the virtual realm is like a massive plain of pristine snow – no matter how lightly we step we leave footprints. What happens to our footprints and the extent to which this matters has been the subject of much angst, debate, speculation, and outrage.

Concerns surrounding personal information being accumulated and stored online fall into several different categories. Those include concerns about criminal or other coercive activity (e.g., identity theft, blackmail, government surveillance and intimidation), behavioral harm (e.g., personal information being used to maximize the addictiveness of platforms like YouTube and Facebook), and concerns around the blurring between our private and public lives (e.g., you cannot count on some regrettable photo meant to remain between friends not popping up to haunt your job interviews years later).

Another major concern, which will be the focus of this paper, is the impact of our online footprints on our economic interactions, which, broadly speaking, include shopping and seeking employment. If you have done any amount of Internet shopping recently it may not come as a surprise that your past shopping history and, more mysteriously, random Google searches, influence the kinds of products that get pitched to you (e.g., on Amazon's homepage, etc.). You may have also noticed that the ads that pop up for you are mysteriously linked to previous online activity (just for fun, you can google fishing-related questions several times in a row and see if you will not start getting ads for fishing gear). Less apparent, but no less troubling, is the fact that the same goods and services may be priced differently based on the online footprints you have left. And even more problematic is the fact that personal information is being accumulated and transformed, through algorithms unknown to the consumer, into output that dictates what jobs and services get pitched to which people. That is, some people may not see, and thus may not get to apply to jobs that some hidden algorithm determined are not likely to be relevant to them based on what their online projection looks like.

The problem as we see it is that this gargantuan information-harvesting apparatus, this expansive web of complex computer algorithms, is largely invisible to the individual and impossible to control or influence with the tools that are available to consumers today. In fact, this apparatus is so insidious that it is often experienced simply as online karma. Meanwhile, on the other end of the economic interaction exist companies that rely on (and pay good money for) the data this apparatus collects, inaccurate as it sometimes is, to enhance their sales and their profit margins.

Current attempts at dealing with this issue have largely been through laws and regulations to improve the security and integrity of online data, and by trying to give consumers a degree of control of the data that is being collected about them (EU's GDPR law, and California's CCPA laws are such recent examples). Some even believe that the problem is overstated and that the market will take care of itself. We would like to claim that the problem is not overstated, is quite serious and will only get worse as computer algorithms become more powerful, and more data about individuals is collected by Internet-connected devices, thus enabling thicker and more expansive

digital portfolios. As for laws and regulations – these are only as valuable as they are enforceable, and we already have multiple regulations that, being poorly fitted to the digital age, are great in theory but toothless in practice. For an example, regulations exist that prohibit various forms of gender discrimination but an individual woman who gets ads for lower paying jobs on average than a male counterpart simply has no way of detecting, let alone addressing this disparity.

We are seeing a rapid increase in the ability of firms to collect data about individuals and to automatically (and thus cheaply) use this data to interact with consumers in ways that benefits the firm. This increase is not matched by a similar increase in individuals' capabilities, making individuals vulnerable for exploitation. We propose that the only sustainable solution that would stop this gap from growing is to arm the consumer with a digital advocate – an equally powerful collection of data and algorithms that would work entirely in the best interest of the individual. To understand what we mean about needing an advocate, perhaps it would be useful to take a detour into a real-world (i.e. non-digital) analogy.

The Car Dealership Analogy

Think back to the last time you bought a car or were involved in shopping for one. Think back to the experience of a car dealership – the giddy smell, the just-scrubbed feeling, the faux chic atmosphere (or maybe it was actually chic – depending on the type of car you were shopping for). And now imagine the person who greeted you there – that impenetrable alluring smile – a firewall behind which lay hidden the information you really needed to know. Information such as, which model is really worth it and which one they are pushing because it gives them higher margin. Or, how low they'd be willing to go on the price if you really pressed them. And if you are like most people (those of you who are true car experts – humor me here), you can probably recall as well the uncomfortable feeling of knowing that in one way or another you are outmatched – that this lovely individual, who on a personal level may really be wishing you well, is actually sizing you up to see how good a deal they can get out of you. Behind that smile, a calculation is always running: How old is this person? Do they seem like they know anything about engines? They are wearing expensive clothes, which means I can try for a higher initial price. They have a small child and, based on past experience, that means I can probably upsell them on safety features. They've come back to look at the same car for a third time. This means they are getting attached to it and will probably be willing to pay more for it. And so, on and on it goes.

If you are like most people, knowing that this kind of calculus occurs makes you feel at least a little uneasy. Even more so because while you can attempt to control what you say, there is really nothing you can do to control how old you look, your accent, your class identifiers and other cues that, statistically speaking, can be used to glean information about you that will ultimately decide what kind of deal you get.

Reflecting on this might inspire some anger (especially if you have vehicular regrets), but the problem with this system is not one of built-in malice. Rather, it is a problem of information imbalance. While the average person shops for a vehicle maybe once or twice in a decade, the

dealership sells dozens of cars every month, so while you are dealing with a situation with an N=1 or 2, they are dealing with a real statistical sample.

Now, imagine how much nicer it might be if you came into a car purchase scenario with your own personal advocate who had access to an information sample about car sales comparable to that available to the dealer. That is, imagine your advocate was privy to all the sales that happened in the past five years in that dealership, the models that sold well and not so well, the average prices offered to people from different groups, etc. On top of that, imagine that this advocate had knowledge of your true ne eds and priorities and could communicate them dispassionately to the dealer (e.g., "they'll be willing to pay a little extra for a better sound system, but please don't try to upsell them on leather seats – they don't need them"). Not only would that be subjectively nicer for the buyer, it might also be objectively more efficient all the way around. Imagine how much time it would save for both parties if the dealer did not have to work so hard on sizing you up and you didn't have to be perpetually on the defensive – time you could use to enjoy your more fairly-priced car and they could use on improving services and selling more cars to a more trusting community.

Back to the digital asymmetry problem

To restate the problem, just as a car dealership uses data from myriad previous customers to dictate their interactions with their current customers, there is a huge apparatus, made up of computer algorithms, that exists to collect data about us online. The purpose of this data collection is ultimately to get us to spend more money and get us to pay the maximum amount of money that we could be induced to pay for the goods and services (to keep our focus, we will ignore for the time being the fact that this data is also used to convince us that we "need" certain goods and services in the first place). Note that we do not claim that this process is malicious by design, just as we do not think that car dealers are malicious. This is simply the natural outcome of an economic interaction where one side has access to lots of data while the other side has access to none. This data collection apparatus is distributed, not regulated, potentially unregulatable, and often impossible to perceive, let alone control, by an individual consumer with the tools available to them today. The information that this apparatus collects, and the inferences it makes are not always accurate or fair, but being that it is currently "the best we got", companies are actively trading in this information between them and with third party procurers, which is a cost that is being passed down to us. As a result, we can assume that we pay an invisible tax every time we engage in commerce online.

To make the impact of this more tangible, think of the last time you shopped online for any item of clothing. Did you really get to choose from the range of products actually available, or were you presented only with higher end brands because somewhere on the Internet there is a record of you buying expensive items in the past? Now imagine that this happens not just with products but also with services and jobs, and not once but multiple times per month.

In a slightly different vein, think of the last time you noticed a weird minor charge on your cellphone bill, or noticed that your insurance plan went up by a few dollars. Now reflect on the

fact that sophisticated automated algorithms exist that allow companies, like insurance and cellphone companies, to calculate, by aggregating data about their customers, the highest surreptitious increase in price the company can tag onto a bill without losing the customer.

So far, we might have made it sound as though having data collected about us is always bad. The really unfortunate thing about all of this is that this is not always the case. Just like there is certain information that we might have liked a car dealer to know if we could fully trust them, there is information that we might like a company to know about us if we could fully trust that it will not be abused. For example, it is not ubiquitously bad for us when products are advertised to us based on what we have searched for previously. The real problem is not that information is being collected, but that this information is being collected behind our backs, outside of our control, and without any input from us. Equally problematic is the fact that while companies can deploy powerful algorithms in their dealings with us, we have no equivalent tools in our arsenals. To illustrate this last point, consider how easy it is for a company to have an electronic system that automatically renews subscriptions to their services and how difficult it can be for any one of us to keep track of all the things we have to unsubscribe from in order to avoid unwanted charges.

The Digital Advocate Solution (and why existing solutions don't work)

We posit that the asymmetry problem described here cannot be solved solely by creating new regulations or tightening existing ones. Existing rules and regulations generally come in two varieties. The first one is prescriptive – of the form "you shall/shall not do such and such" (e.g. "you shall not sell toys that contain lead", or "you shall give customers fair warning before increasing a rate"). This type of regulation actually tends to work quite well when the object of the injunction is easily measured but becomes problematic and easy to bypass when its object is not easy to measure or define. To illustrate, it is very difficult to equivocate about what it means for paint to contain lead. Either there is lead in that paint or there is not. Toy companies may, and have, tried to play around with how much lead is harmful, or tried to mislead consumers by hiding the fact that there is lead in their toys, but there is no hiding from the fact that there exist laboratories that can definitively answer the question: "is there lead in this toy?" and thus "did this company break the rule?". Now, try the same exercise with the "fair warning" rule. Suppose you signed up for an online service (say a streaming website) that you later learned automatically increased your rate by 50% after the first three months of use. We can all agree that they will have broken the rule if they never mentioned this rate increase anywhere at all. But what about if they mentioned it on page 15 of their user agreement? What if they mentioned it on page 15 of an agreement that only opens when you click a button that most people would never click? What if they only mentioned it on page 15 of a document that they sent in an attachment in a follow-up email that contained a user satisfaction survey in the body of the email? Did they break the rule in those situations? They certainly broke the spirit of the rule, but there is a lot of ground for argument about whether they broke the letter of the rule, so as you can see, prescriptive regulations can be ambiguous and easy to bypass.

The second kind of regulation is of the form "if X happens to you, you can sue the perpetrator". For example, if you sustain bodily injury while shopping and you can demonstrate that this was due to negligence on the store's part (e.g. a wet floor, an exposed live wire, etc.), you can sue the store. This form of regulation, while definitely necessary, tends to be very reactive and based on hindsight being 20/20. It is usually not until something happens to someone that a new regulation form. As such, this form of regulation can work in an arena where things stay relatively constant, but is too slow for the digital realm, where things change daily. But perhaps the biggest problem with this form of regulation is that it presupposes a specific, identifiable entity that can be seen to bear responsibility for whatever transpired. If a specific responsible entity cannot be identified it is really not clear who is to be sued even if we could agree that a transgression occurred. Now, suppose that an exploitative event happened to you while shopping on a site that contracts with three other entities to deliver its services (all, by the way, unknown to the consumer) and sells items from independent third-party vendors. Even if you somehow managed to prove that an illegal even (say discrimination) happened to you, you can probably see that untangling this complex web of deniability, if possible, is not an easy task.

As far as we can see, it is a daunting and perhaps impossible task to get rid of bad practices on the Internet simply by regulating them out of existence. Rules and regulations are simply not quick or agile enough for that. What we envision is a solution that will approach this problem entirely differently. The solution we propose will shift the balance of power between individual consumers and companies in such as a way as to make individual consumers more equally matched with the organizations they transact with. Similar to having a personal advocate at the car dealership, we envision an algorithm equal in power to those used by large entities working entirely in the individual's best interest and accountable only to the individual who retained it. Such an algorithm will have several important features.

Firstly, it will be able to gain access to and take charge of the information that is collected about us online. This will give us the power to monitor and influence this information and also to trade in that information ourselves, which is a vast improvement upon the current situation, where shady entities trade in information about us behind our backs. Secondly, it will have access to information about the environment in a scope that a single individual simply cannot possess. For instance, in a shopping scenario, it will have access to statistics about all similar transactions within the past five years and have the ability to share information with other algorithms representing other individuals. As a result, it will be able to alert us if the range of products we are seeing is truncated and even tell us what part of the full range is being excluded. All of this might help us get the sneakers that are best for us chosen from a range of hundreds of products, as opposed to getting the sneakers that are being pushed on us from a range of ten products. The algorithm will also be able to alert us if the price being offered to us is statistically different from the average price offered to others. It will be able to communicate to the seller on our behalf and let them know that we know what the fair price is and will pay no more than that. Now stretch your imagination even further. What if you actually would be willing to pay a slightly

higher price in exchange for expedited shipping, or a festive package, or some other perk that is relevant for you. Such nuances of haggling, formerly ubiquitously practiced, are very difficult to imagine while shopping online – the sellers are currently simply too powerful to haggle with. With a digital advocate algorithm haggling could make its way back into commerce – your algorithm could haggle with the seller on an equal footing.

Even in terms of more traditional regulation, such digital advocates appear to be sorely missing. Suppose it were illegal for car dealers to routinely offer older buyers higher prices on the same car. How would you detect such a behavior? The best solution would be to send mystery shoppers to conduct a randomized controlled trial – which is a tall order even for a digital advocate. Short of that, the next best thing would be to collect a representative sample of transactions and analyze it for price gouging, and for that a digital advocate would be perfect.

In the realm of digital account management, you digital advocate will be able to monitor your accounts and automatically notify your insurance company that it noticed a 7% rate increase in the absence of an industry-wide price increase or any precipitating event, and that unless they explain themselves you will switch providers. Your digital advocate could also take full advantage of all those wonderful consumer protection clauses that many companies and states have on paper but in practice are rarely used. People are not very likely to spend an hour to dispute a ten-dollar charge, but repeated thousands of times over someone's lifetime this can add up to a lot of money. You can probably see how a digital advocate could be very useful here. Undaunted by clutter and infinitely attentive to detail it could finally get us those rebates, money back guarantees and reversals of spurious charges that we currently let slide. Better yet, if it became common knowledge that bait-and-switch is not a viable business model, the prevalence of such practices would drop dramatically.

Finally, and most importantly, our economic 'selves' are steadily migrating online. As our online digital footprints become more complex and more all-encompassing, it gets exponentially harder to have control over the effect these footprints have on our lives, or even 'just' economic aspects of our lives. Even when favorable laws, such as the right to download one's data under GDPR, are in place, it is often difficult to actually get the companies to comply. Even if they comply, the brave souls who tried (mostly privacy connoisseurs and journalists) found that it is difficult for an expert to make sense of these data dumps, let alone a lay person. This data represents tremendous economic value to the company that produced it, but putting this data to work *for us* outside of control of its control seems like science fiction in today's digital environments. These data were produced and collected by algorithms, and it will take another algorithm to make use of it. This is our only hope for keeping a degree of control over our economic selves.

Why does a digital advocate not exist yet, and what would it take to develop it

You might ask – if a digital advocate is such a great idea, why has it not been developed already. Well – one answer, valid but perhaps not very interesting, is that the digital world is still

new and it has been growing and evolving at breakneck speed, so perhaps the digital advocate does not exist yet simply because technologies and business models to enable it have not evolved yet. This is not very convincing, given that the Internet has been around for over 25 years, and *business* analogues of the digital advocate (digital customer relationship management, business analytics, business digital marketing management) have existed for as long (and sometimes longer).

One glaring issue is an issue of trust. As you read through the last few pages this might have bothered you – if this digital advocate is to work properly, we would need to share a whole lot of very personal information with it. Why would we want to entrust our already vulnerable online selves to another powerful algorithm? How do we know it will not betray us? In this context people usually worry about two separate issues.

The first of these (and this tends to scare people quite a bit) is the worry that this digital advocate might expose us to privacy and security violations – by being hacked or taken over by bad actors. While this is an important issue, it is unlikely to be the issue that had stopped digital advocates from being implemented. For one, it is unlikely that such an advocate would have more information about an individual than Google, Internet service providers such as Verizon or for that matter some of the back-end Internet trackers already have thorough Internet search and browsing histories. In fact, there are plenty of tools that serve as information aggregators (such as financial or medical insurance claims) which already have access to extremely sensitive information, but usually stop short of full advocacy, instead using a combination of advertising and market research as their business model.

The second issue is that of economic trust: will such an advocate act in our best interest (as its name suggests), or in someone else's? The problem of economic trust is much more significant. On the matter of privacy and security the advocate's and the client's incentives are fully aligned: both would like to prevent unauthorized breaches. On the other hand, there is an inherent economic conflict of interest between an advocate (digital or human) and its client. An advocate who has access to a client's private information, can use this access to its own benefit. This can be outright harmful to the client (for example steering clients with a higher willingness to pay towards higher-priced products, then pocketing a commission), but often has a more neutral-appearing flavor (steering a client towards one of several hard-to-compare options, then collecting a commission). The latter is the preferred business model of most consumer-facing platforms and tools.

The problem with such a business model is that it leads to a limited utility for the consumer. An advocate interfacing with a consumer (who is limited in many ways as we discussed and can't readily assess the benefit of the advocate) on one end, and with companies equipped with sophisticated algorithms and marketing cash on the other, will end up gravitating towards the latter. In a world where it is impossible to commit to "not sell out" the consumer, or where it is difficult to get a competitive advantage based on such a commitment, attempts of constructing a digital advocate will (and have) end up with a product that trades off some utility for the consumer (such as search results, convenient information access, or coupons) in exchange

for a hidden cost (a hidden commission on a sale, users' activity data sold to third parties). This just leads to another layer of technology on which the individual user depends, but which is not economically accountable to her.

A closely related issue is the issue of cost. If an advocate is not allowed to receive commission based on products it helps sell, then the consumer will have to pay the advocate directly. Facebook costs well under $100/user/year to run. At scale, a digital advocate should not cost more than that. Considering the impact our digital interactions have on our lives (economically and beyond), this cost appears almost trivial. The abundance of "shopping assistant" tools online, suggest that marketing commissions alone are enough to pay for such an advocate. Thus "sell user out for commissions but bring the commissions to the user in cash" is a potentially viable "free" business model – the reason it hasn't been implemented is that absent strong regulations/norms, this model loses to the "sell user out for commissions, give user portion of commissions, use portion of commissions to make product nicer, pocket the rest" model. In addition, having a digital advocate has the potential of not just being good for an individual user, but for the digital economy as a whole, as it will make it easier to uncover and discourage unethical business practices, while encouraging competition.

Finally, there are technical hurdles. By far the hardest hurdle is establishing and maintaining data links to maintain an accurate digital "portrait" of the user. As we have seen earlier, there is no simple answer to the question "what does the Internet know about you", even if you just narrow it down to the economic dimension. In some cases, a user has the right to request her "dossier" from an Internet company such as Facebook, but such "rights" are of limited practical value, since an algorithmic tool (such as the proposed advocate) is needed to make any use of the data. This problem is exacerbated by two considerations. First, that our digital portrait is distributed among many players, large and small, which may fall under multiple jurisdictions and have different data export policies. Second, what one needs is not merely a snapshot but a continuously updated "live" history. There are many tools that aggregate data over a specific domain, such as messages, or financial account information, or credit monitoring, but all these tools generally work with data that is already readily exposed to the user – capturing data that is currently not being exposed (and is sometimes deliberately obfuscated) is a different matter, which will need to be addressed by regulators.

Concrete steps

To recap, we have argued that there is a growing gap between the ability of companies to deploy algorithms to obtain information about individual users and to extract economic advantages from it, sometimes at the expense of the user, and the user's ability to counter these algorithms. The gap between the algorithmic capabilities of the data-collecting machines and the ability of the individual users keeps increasing, and currently there is no serious force to keep this gap from expanding indefinitely. Regulation in its current format is unlikely to stop it from expanding further, as the economic logic in its favor is too overwhelming. A personal digital advocate equipped with algorithmic power comparable to that of the companies at the other

end of the equation is a potential way to restore the balance between individuals (and, by extension, governments) and companies.

Unlike many technological tools that developed organically as technologies matured, there are two policy hurdles that need to be overcome before they can become a reality. We conclude by briefly discussing those, as we believe addressing them should be a policy priority – perhaps more important than expanding other forms of regulation on existing corporate players.

Creating a regulatory environment of trust. Having digital access to a large amount of behavioral and economic data from a user gives the digital agent a lot of power to steer user behavior. This power can be used to the user's benefit, but it can also be used for the benefit of a third party in ways that are not beneficial or even harmful to the user. Moreover, the power disparity is sufficiently significant that such arrangements, even with an "opt out" or even an "opt in" regime are still likely to end up being exploitative. The only way to prevent such an outcome is to explicitly create a legal structure under which the digital agent is bound to working exclusively in the best interest of the user. Such a regime would be akin to the way relationships with certain professionals are regulated. A lawyer or a psychologist who abuses their position of trust for their own benefit risks promptly losing their license (and probably worse). A client cannot "opt out" of these licensing rules – e.g. by allowing their psychologist to sell information about them to marketers in exchange for reduced fees.

This would probably be best accomplished by a regulatory board, similar to professional boards as in Law and Medicine, which would create standards for what it means for an algorithm to act in the "user's best interest", educate the public about its rights and expectations, and enforce these standards through some form of a licensing regime.

Technical regulation of data import/export. The first instinct of anyone faced with the existence of a dossier that affects them is to access it to correct mistakes, and to influence it to one's benefit. Indeed, many regulations around such dossiers (such as credit histories, medical records, education records) enforce mechanisms allowing their subjects to access records about them. Existing pushes in the domain of Internet privacy, such as GDPR, make the ability to access such data a priority. The next logical step in this process (currently missing) is to require all data to be available in a standardized machine-readable format (which the advocate will be able to access). While currently the company has 30 days to respond to a data request, there should be a regime in which most recent data is made accessible to the advocate in regular intervals (daily or weekly).

Similarly, if one views online privacy and fairness regulations as 'Bills of Rights', such rights should come with a machine-friendly implementation requirement: if I am allowed to opt-out of tracking by a given ad network, there should be a link my digital advocate can access to make the request, without needing a human in the loop (beyond giving the general instruction to my advocate to opt-out).