

12 Marginal People on Marginal Time

How can we begin to assert and defend our freedom of attention? One thing is clear: it would be a sad reimposition of the same technocratic impulse that gave us the attention economy in the first place if we were to assume that there exists a prescribable basket of “solutions” which, if we could only apply them faithfully, might lead us out of this crisis. There are no maps here, only compasses. There are no three-step templates for revolutions.

We can, however, describe the broad outline of our goal: it’s to bring the technologies of our attention onto our side. This means aligning their goals and values with our own. It means creating an environment of incentives for design that leads to the creation of technologies that are aligned with our interests from the outset. It means being clear about what we want our technologies to do for us, as well as expecting that they be clear about what they’re designed to do for us. It means expecting our technologies to proceed from a place of understanding about our own views of who we are, what we’re doing, and where we’re going. It means expecting our technologies and their designers to give attention to, to care about, the right things. If we move in the right direction, then our fundamental understanding of what technology is *for*, as the philosopher Charles Taylor has put it, “will of itself be limited and enframed by an ethic of caring.”¹

Drawing on this broad view of the goal, we can start to identify some vectors of rebellion against our present attentional serfdom. I don’t claim to have all, or even a representative set, of the answers here. Nor is it clear to me whether an accumulation of incremental improvements will be sufficient to change the system; it may be that some more fundamental reboot of it is necessary. Also, I won’t spend much time here talking about who in society bears responsibility for

putting each form of attentional rebellion into place: that will vary widely between issues and contexts, and in many cases those answers aren't even clear yet.

Prior to any task of systemic reform, however, there's one extremely pressing question that deserves as much of our attention as we're able to give it. That question is whether there exists a "point of no return" for human attention (in the deep sense of the term as I have used it here) in the face of this adversarial design. That is to say, is there a point at which our essential capacities for life navigation might be so undermined that we would become unable to regain them and bootstrap ourselves back into a place of general competence? In other words, is there a "minimum viable mind" we should take great pains to preserve? If so, it seems a task of the highest priority to understand what that point would be, so that we can ensure we do not cross it. In conceiving of such a threshold – that is, of the minimally necessary capacities worth protecting – we may find a fitting precedent in what Roman law called the "benefit of competence," or *beneficium competentiae*. In Rome, when a debtor became insolvent and couldn't pay his debts, there was a portion of his belongings that couldn't be taken from him in lieu of payment: property such as his tools, his personal effects, and other items necessary to enable a minimally acceptable standard of living, and potentially even to bootstrap himself back into a position of thriving. This privileged property that couldn't be confiscated was called his "benefit of competence." Absent the "benefit of competence," a Roman debtor might have found himself ruined, financially destitute. In the same way, if there is a "point of no return" for human attention, a "minimum viable mind," then absent a "benefit of competence" we could also find ourselves ruined, attentionally destitute. And we are not even debtors: we are serfs in the attentional fields of our digital technologies. They are in *our* debt. And they owe us, at absolute minimum, the benefit of competence.

There are a great number of interventions that could help move the attention economy in the right direction. Any one could fill a whole book. However, four particularly important types I'll briefly

discuss here are: (a) rethinking the nature and purpose of advertising, (b) conceptual and linguistic reengineering, (c) changing the upstream determinants of design, and (d) advancing mechanisms for accountability, transparency, and measurement.

If there's one necessary condition for meaningful reform of the attention economy, it's the reassessment of the nature and purpose of advertising. It's certainly no panacea, as advertising isn't the only incentive driving the competition for user attention. It is, however, by far the largest and most deeply ingrained one.

What is advertising *for* in a world of information abundance? As I wrote earlier, the justification for advertising has always been given on the basis of its informational merits, and it has historically functioned within a given medium as the *exception to the rule* of information delivery: for example, a commercial break on television or a billboard on the side of the road. However, in digital media, advertising now *is* the rule: it has moved from "underwriting" the content and design goals to "overwriting" them. Ultimately, we have no conception of what advertising is *for* anymore because we have no coherent definition of what advertising *is* anymore.²

As a society, we ought to use this state of definitional confusion as the opportunity to help advertising resolve its existential crisis, and to ask what we ultimately want advertising to do for us. We must be particularly vigilant here not to let precedent serve as justification. As Thomas Paine wrote in *Common Sense*, "a long habit of not thinking a thing wrong, gives it a superficial appearance of being right."³ The presence of a series of organizations dedicated to a task can in no sense be justification for that task. (See, e.g., the tobacco industry.) What forms of attitudinal and behavioral manipulation shall we consider to be acceptable business models? On what basis do we regard the wholesale capture and exploitation of human attention as a natural or desirable thing? To what standards ought we hold the mechanisms of commercial persuasion, knowing full well that they will inevitably be used for political persuasion as well?

A reevaluation of advertising's *raison d'être* must necessarily occur in synchrony with the resuscitation of serious advertising ethics. Advertising ethics has never really guided or restrained the practice of advertising in any meaningful way: it's been a sleepy, tokenistic undertaking. Why has this been so? In short, because advertisers have found ethics threatening, and ethicists have found advertising boring. (I know, because I have been both.)

In advertising parlance, the phrase "remnant inventory" refers to a publisher's unpurchased ad placements, that is, the ad slots of *de minimis* value left over after advertisers have bought all the slots they wanted to buy. In order to fill remnant inventory, publishers sell it at extremely low prices and/or in bulk. One way of viewing the field of advertising ethics is as the "remnant inventory" in the intellectual worlds of advertisers and ethicists alike.

This general disinterest in advertising ethics is doubly surprising in light of the verve that characterized voices critical of the emerging persuasion industry in the early to mid twentieth century. Notably, several of the most prominent early critical voices were veterans of the advertising industry. In 1928, brand advertising luminary Theodore MacManus published an article in the *Atlantic Monthly* titled "The Nadir of Nothingness" that explained his change of heart about the practice of advertising: it had, he felt, "mistaken the surface silliness for the sane solid substance of an averagely decent human nature."⁴ A few years later, in 1934, James Rorty, who had previously worked for the McCann and BBDO advertising agencies, penned a missive titled *Our Master's Voice: Advertising*, in which he likewise expressed a sense of dread that advertising was increasingly violating some fundamental human interest:

[Advertising] is never silent, it drowns out all other voices, and it suffers no rebuke, for is it not the voice of America? . . . It has taught us how to live, what to be afraid of, how to be beautiful, how to be loved, how to be envied, how to be successful . . . Is it any wonder that the American population tends increasingly to speak, think, feel in terms

of this jabberwocky? That the stimuli of art, science, religion are progressively expelled to the periphery of American life to become marginal values, cultivated by marginal people on marginal time?⁵

The prose of these early advertising critics has a certain tone, well embodied by this passage, that for our twenty-first-century ears is nearly impossible to ignore. It's a sort of pouring out of oneself, an expression of disbelief and even offense at the perceived aesthetic and moral violations of advertising, and it's further tinged by a plaintive, interrogative style that reminds us of other Depression-era writers (James Agee in particular comes to mind). But it reminds me of Diogenes, too: when he said he thought the most beautiful thing in the world was "freedom of speech," the Greek word he used was *parthesia*, which doesn't just mean "saying whatever you want" – it also means speaking boldly, saying it all, "spilling the beans," pouring out the truth that's inside you. That's the sense I get from these early critics of advertising. In addition, there's a fundamental optimism in the mere fact that serious criticism is being leveled at advertising's existential foundations *at all*. Indeed, reading Rorty today requires a conscious effort not to project our own rear-view cynicism on to him.

While perhaps less poetic, later critics of advertising were able to more cleanly circumscribe the boundaries of their criticism. One domain in which neater distinctions emerged was the logistics of advertising: as the industry matured, it advanced in its language and processes. Another domain that soon afforded more precise language was that of psychology. Consider Vance Packard, for instance, whose critique of advertising, *The Hidden Persuaders* (1957), had the benefit of drawing on two decades of advances in psychology research after Rorty. Packard writes: "The most serious offense many of the depth manipulators commit, it seems to me, is that they try to invade the privacy of our minds. It is this right to privacy in our minds – privacy to be either rational or irrational – that I believe we must strive to protect."⁶

Packard and Rorty are frequently cited in the same neighborhood in discussions of early advertising criticism. In fact, the frequency with

which they are jointly invoked in contemporary advertising ethics research invites curiosity. Often, it seems as though this is the case not so much for the content of their criticisms, nor for their antecedence, but for their tone: as though to suggest that, if someone were to express today the same degree of unironic concern about the foundational aims of the advertising enterprise as they did, and to do so with as much conviction, it would be too embarrassing, quaint, and optimistic to take seriously. Perhaps Rorty and Packard are also favored for their perceived hyperbolizing, which makes their criticism easier to dismiss. Finally, it seems to me that anchoring discussions about advertising's fundamental ethical acceptability in the distant past may have a rhetorical value for those who seek to preserve the status quo; in other words, it may serve to imply that any ethical questions about advertising's fundamental acceptability have long been settled.

My intuition is that the right answers here will involve moving advertising away from *attention* and towards *intention*. That is to say, in the desirable scenario advertising would not seek to capture and exploit our mere *attention*, but rather support our *intentions*, that is, advance the pursuit of our reflectively endorsed tasks and goals.

Of course, we will not reassess, much less reform, advertising overnight. Until then, we must staunchly defend, and indeed enhance, people's ability to decline the harvesting of their attention. Right now, the practice currently called "ad blocking" is one of the only ways people have to cast a vote against the attention economy. It's one of the few tools users have if they want to push back against the perverse design logic that has cannibalized the soul of the web. Some will object and say that ad blocking is "stealing," but this is nonsense: it's no more stealing than walking out of the room when the television commercials come on. Others may say it's not prudent to escalate the "arms race" – but it would be fantastic if there were anything remotely resembling an advertising arms race going on. What we have instead is, on one side, an entire industry spending billions of dollars trying to capture your attention using the most sophisticated computers in the world, and on the other side . . . your attention. This is

more akin to a soldier seeing an army of thousands of tanks and guns advance upon him, and running into a bunker for refuge. It's not an arms race – it's a quest for attentional survival.

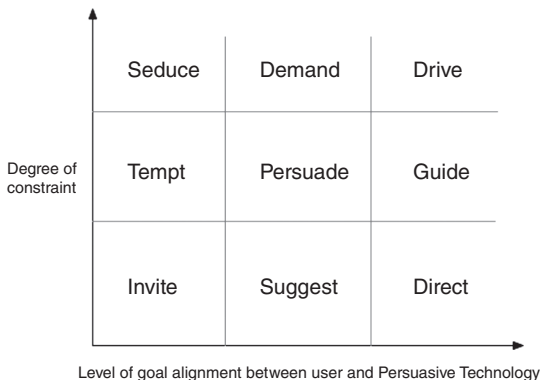
The right of users to exercise and protect their freedom of attention by blocking any advertising they wish should be absolutely defended. In fact, given the moral and political crisis of the digital attention economy, the relevant ethical question here is not "Is it okay to block ads?" but rather, "Is it a moral obligation?" This is a question for companies, too. Makers of digital technology hardware and software ought to think long and hard about their obligations to their users. I would challenge them to come up with any good reasons why they shouldn't ship their products with ad blocking enabled by default. Aggressive computational persuasion should be opt-in, not opt-out. The default setting should be one of having control over one's own attention.

Another important bundle of work involves reengineering the language and concepts of persuasive design. This is necessary not only for talking clearly about the problem, but also for advancing philosophical and ethical work in this area. Deepening the language of "attention" and "distraction" to cover more of the human will has been part of my task here. Concepts from neuroethics may also be of help in advancing the ethics of attention, especially in describing the problem and the nature of its harms, as in, for example, the concepts of "brain privacy" or "cognitive liberty."⁷

For companies, a key piece of this task involves reengineering the way we talk about users. Designers and marketers routinely use terms like "eyeballs," "funnels," "targeting," and other words that are perhaps not as humanized as they ought to be. The necessary corrective is to find more human words for human beings. To put a design spin on Wittgenstein's quote from earlier, we might say that the limits of our language mean the limits of our empathy for users.

Regarding the language of "persuasion" itself, there is a great deal of clarification, as well as defragmentation across specific contexts of persuasion, that needs to occur. For example, we could map

the language of “persuasive” technologies according to certain ethically salient criteria, as seen in the figure below, where the Y axis indicates the level of constraint the design places on the user and the X axis indicates the degree of alignment between the user’s goals and the technology’s goals. Using this framework, then, we could describe a technology with a low level of goal alignment and a high degree of constraint as a “Seductive Technology” – for example, an addictive game that a user wants to stop playing, and afterward regrets having spent time on. However, if its degree of constraint were very low, we could instead call it an “Invitational Technology.” Similarly, a technology that imposes a low degree of constraint on the user and is highly aligned with their goals, such as a GPS device, would be a “Directive Technology.” As its constraints on the user increase, it would become a “Guidance Technology” (e.g. a car’s assisted-parking or autopilot features) and at even higher levels a “Driving Technology” (e.g. a fully autonomous vehicle). This particular framework is an initial, rough example for demonstrating what I mean, but it illustrates some of the ways such a project of linguistic and conceptual defragmentation could go.



Clarifying the language of persuasion will have the added benefit of ensuring that we don’t implicitly anchor the design ethics of

attention and persuasion in questions of *addiction*. It's understandable why discussion about these issues has already seized on addiction as a core problem: the fundamental challenge we experience in a world of information abundance is a challenge of self-control, and the petty design habits of the attention economy often target our reward system, as I described in Chapter 4.

But there are problems with giving too much focus to the question of addiction. For one, there's a strict clinical threshold for addiction, but then there's also the colloquial use of the term, as shorthand for "I use this technology more than I want to." Without clear definitions, it's easy for people to talk past one another. In addition, if we give too much focus to addiction there's the risk that it could implicitly become a default threshold used to determine whether a design is morally problematic or not. But there are many ways a technology can be ethically problematic; addiction is just one. Even designs that create merely compulsive, rather than "addictive," behaviors can still pose serious ethical problems. We need to be especially vigilant about this sort of ethical scope creep in deployments of the concept of addiction because there are incentives for companies and designers to lean into it: not only does this set the ethical threshold at a high as well as vague level, but it also serves to deflect ethical attention away from deeper ethical questions about goal and value misalignments between the user and the design. In other words, keeping the conversation focused on questions of addiction serves as a convenient distraction from deeper questions about a design's fundamental purpose.

Interventions with the highest leverage would likely involve changing the upstream determinants of design. This could come from, for instance, the development and adoption of alternate corporate structures that give companies the freedom to balance their financial goals with social good goals, and then offer incentives for companies to adopt these corporate structures. (For instance, Kickstarter recently transitioned to become a "benefit corporation,"

or B-corp. The writer and Columbia professor Tim Wu has recently called in the *New York Times* for Facebook to do the same.)⁸ Similarly, investors could create a funding environment that disincentivizes startup companies from pursuing business models that involve the mere capture and exploitation of user attention. In addition, companies could be expected (or compelled, if necessary) to give users a choice about how to “pay” for content online – that is, with their money or with their attention.

Many of these upstream determinants of design may be addressed by changes in the policy environment. Policymakers have a crucial role to play in responding to the crisis of the digital attention economy. To be sure, they have several headwinds working against them: the internet’s global nature means local policies can only reach so far, and the rapid pace of technological change tends to result in reactive, rather than proactive, policymaking. But one of the strongest headwinds for policy is the persistence of *informational*, rather than *attentional*, emphases. Most digital media policy still arises out of assumptions that fail to sufficiently account for Herbert Simon’s observation about how information abundance produces attention scarcity. Suggestions that platforms be required to tag “fake news,” for example, would be futile, an endless game of epistemic whack-a-mole. Initial research has already indicated as much.⁹ Similarly, in the European Union, website owners must obtain consent from each user whose browsing behavior they wish to measure via the use of tracking “cookies.” This law is intended to protect user privacy and increase transparency of data collection, both of which are laudable aims when it comes to the ethics of *information* management. However, from the perspective of *attention* management, the law burdens users with, say, thirty more decisions per day (assuming they access thirty websites per day) about whether or not to consent to being “cookied” by a site they may have never visited before, and therefore don’t know whether or not they can trust. This amounts to a nontrivial strain on their cognitive load that far outweighs any benefit of giving their “consent” to have their browsing behavior measured. I place the word

“consent” in quotes here because what inevitably happens is that the “cookie consent” notifications that websites show to users simply become designed to maximize compliance: website owners simply treat the request for “consent” as one more persuasive interaction, and deploy the same methods of measurement and experimentation they use to optimize their advertising-oriented design in order to manufacture users’ consent.

However, governmental bodies are uniquely positioned to host conversations about the ways new technological affordances relate to the moral and political underpinnings of society, as well as to advance existential questions about the nature and purpose of societal institutions. And, importantly, they are equipped to foster these conversations in a context that can, in principle, inform and catalyze corrective action. We can find some reasons to be at least cautiously optimistic in precedents for legal protection of attention enacted in predigital media. Consider, for instance, anti-spam legislation and “do not call” registries, which aim to forestall unwanted intrusions into people’s private spaces. While protections of this nature generally seek to protect “attention” in the narrow sense – in other words, to mitigate annoyance or momentary distractions – they can nonetheless serve as doorways to protecting the deeper forms of “attention” that I have discussed here.

What can policy do in the near term that would be high-leverage? Develop and enforce regulations and/or standards about the transparency of persuasive design goals in digital media. Set standards for the measurement of certain sorts of attentional harm – that is, quantify their “pollution” of the *inner* environment – and require that digital media companies measure their effects on these metrics and report on them periodically to the public. Perhaps even charge companies something like carbon offsets if they go over a certain amount – we might call them “attention offsets.” Also worth exploring are possibilities for digital media platforms that would play a role analogous to the role public broadcasting has played in television and radio.

Advancing accountability, transparency, and measurement in design is also key. For one, having transparency of persuasive design goals is essential for verifying that our trust in the creators of our technologies is well placed. So far, we've largely demanded transparency about the ways technologies manage our *information*, and comparatively less about the ways they manage our *attention*. This has foregrounded issues such as user privacy and consent, issues which, while important, have distracted us from demanding transparency about the design logic – the ultimate *why* – that drives the products and services we use. The practical implication of this is that we've had minimal and shaky bases for trust. "Whatever man you meet," advised the Roman emperor Marcus Aurelius in his *Meditations*, "say to yourself at once: 'what are the principles this man entertains as goods and ills?'"¹⁰ This is good advice not only upon encountering persuasive people, but persuasive technologies as well. What is Facebook's persuasive goal for me? On what basis does YouTube suggest that I watch one video and not another? What metric does Twitter aim to maximize with my time use? Why *did* Amazon build Alexa, after all? Do the goals my trusted systems have for me align with the goals I have for myself? There's nothing wrong with trusting the people behind our technologies, nor do we need perfect knowledge of their motivations to justifiably do so. Trust always involves taking *some* risk. Rather, our aim should be to find a way, as the Russian maxim says, to "trust, but verify."

Equipping designers, engineers, and businesspeople with effective "commitment devices" may also be of use. One common example is that of professional oaths. The oath occupies a unique place in contemporary society: it's weightier than a promise, more universal than a pledge, and more individualized than a creed. Oaths express and remind us of common ethical standards, provide opportunities for making public commitments to particular values, and enable accountability for action. Among the oaths that are not legally binding, the best known is probably the Hippocratic Oath, some version of which is commonly recited by doctors when they graduate from

medical school. Karl Popper (in 1970)¹¹ and Joseph Rotblat (in 1995),¹² among others, have proposed similar oaths for practitioners of science and engineering, and in recent years proposals for oaths specific to digital technology design have emerged as well.¹³ So far, none of these oaths have enjoyed broad uptake. The reasons for this likely include the voluntary nature of such oaths, as well as the inherent challenge of agreeing on and articulating common values in pluralistic societies. But the more significant headwinds here may originate in the decontextualized ways in which these proposals have been made. If a commitment device is to be adopted by a group, it must carry meaning for that group. If that meaning doesn't include some sort of *social* meaning, then achieving adoption of the commitment device is likely to be extremely challenging. Most oaths in wide use today depend on some social structure below the level of the profession as a whole to provide this social meaning. For instance, mere value alignment among doctors about the life-saving goals of medicine would not suffice to achieve continued, widespread recitations of the Hippocratic Oath. The essential infrastructure for this habit lies in the social structures and traditions of educational institutions, especially their graduation ceremonies. Without a similar social infrastructure to enable and perpetuate use of a "Designer's Oath," significant uptake seems doubtful.

It could be argued that a "Designer's Oath" is a project in search of a need, that none yet exists because it would bring no new value. Indeed, other professions and practices seem to have gotten along perfectly fine without common oaths to bind or guide them. There is no "Teacher's Oath," for example; no "Fireman's Oath," no "Carpenter's Oath." It could be suggested that "design" is a level of abstraction too broad for such an oath because different domains of design, whether architecture or software engineering or advertising, face different challenges and may prioritize different values. In technology design, the closest analogue to a widely adopted "Designer's Oath" we have seen is probably the voluntary ethical commitments that have been made at the organizational level, such as company

mottos, slogans, or mission statements. For example, in Google's informal motto, "do no evil," we can hear echoes of that Hippocratic maxim, *primum non nocere* ("first, do no harm").¹⁴

But *primum non nocere* does not, in fact, appear in any version of the Hippocratic Oath. The widespread belief otherwise provides us with an important signal about the perceived versus the actual value of oaths in general. A significant portion of their value comes not from their content but from their mere existence: from the societal recognition that a particular practice or profession is oath-worthy, that it has a significant impact on people's lives such that *some* explicit ethical standard has been articulated to which conduct within the field can be held.

Assuming we could address these wider challenges that limit the uptake of a "Designer's Oath" within society, what form should such an oath take? In this space, I can only gesture toward a few of the main questions – let alone arrive at any clear answers. One of the key questions is how explicitly such an oath should draw on the example of the Hippocratic Oath. In my view, the precedent seems appropriate to the extent that using the metaphor of medicine to talk about design can help people better understand the seriousness of design. Comparing design to medicine is a useful way of conveying the depth of what is ultimately at stake. Medicine is also an appropriate metaphor because, like design, it's a profession rather than an organization or institution, which makes it an appropriate level of society at which to draw a comparison.

However, one limitation of drawing on medicine as a rough guide to this terrain pertains to the logistics of when and where (and by whom) a "Designer's Oath" would be taken. Medical training is highly systematized, and provides an organizational context for taking such an oath. A technology designer, by contrast, may have never had *any* formal design education – and even those who have, may have never taken a design ethics class. Even for those who *do* take design ethics classes (which are often electives), there is unlikely to be a moment in them when, as in a graduation ceremony, it would

not feel extremely awkward to take an oath. Of course, this assumes that an educational setting is the appropriate context for such an oath to begin with. Should we instead look to companies to lead the way? If so, this would raise the further question about *who* should be expected, and not expected, to take the oath (e.g. front-end vs. back-end designers, hands-on designers vs. design researchers, senior vs. junior designers, etc.). Finally, there's also the question of *how* such an oath should be written, especially in the digital age. Should it be a "wiki"-style oath, the product of numerous contributors' input and discussion? Or is such a "crowd-sourced" approach, while an appropriate way to converge on the provisional truth of a *fact* (as in Wikipedia), an undesirable way to develop a clear-minded expression of a moral *ideal*? In any event, we should expect that any "Designer's Oath" receiving wide adoption would continually be iterated and adapted in response to local contexts and new advances in ethical thought, as has been the case with the Hippocratic Oath over many centuries.

As regards the *substance* of a "Designer's Oath" – an initial "alpha" version that can serve as a "minimum viable product" to build upon – I suggest that a good approach would look something like the following (albeit far more poetic and memorable than this):

As someone who shapes the lives of others, I promise to:

- Care** genuinely about their success;
- Understand** their intentions, goals, and values as completely as possible;
- Align** my projects and actions with their intentions, goals, and values;
- Respect** their dignity, attention, and freedom, and never use their own weaknesses against them;
- Measure** the full effect of my projects on their lives, and not just those effects that are important to me;
- Communicate** clearly, honestly, and frequently my intentions and methods; and
- Promote** their ability to direct their own lives by encouraging reflection on their own values, goals, and intentions.

I won't attempt here to justify each element I've included in this "alpha" version of the oath, but will only note that: (a) it assumes a patient-centered, rather than an agent-centered, perspective; (b) in keeping with the theme of this inquiry, it emphasizes ethical questions related to the management of attention (broadly construed) rather than the management of information; (c) it explicitly disallows design that is consciously adversarial in nature (i.e. having aims contrary to those of the user), which includes a great deal of design currently operative in the attention economy; (d) it goes beyond questions of respect or dignity to include an expectation of *care* on the part of the designer; and (e) it views measurement as a key way of operationalizing that care in the context of digital technology design (as I will further discuss below).

Measurement is also key. In general, our goal in advancing measurement should be to measure what we value, rather than valuing what we already measure. Ethical discussions about digital advertising often assume that limiting user measurement is axiomatically desirable due to considerations such as privacy or data protection. These are indeed important ethical considerations, and if we conceive of the user-technology interaction in informational terms then such conclusions may very well follow. Yet if we take an attention-centric perspective, as I have described above, there are ways in which limiting user measurement may complicate the ethics of a situation, and possibly even actively hinder it.

Greater measurement (of the right things) is in principle a good thing. Measurement is the primary means designers and advertisers have of attending to specific users, and as such it can serve as the ground on which conversations, and if necessary interventions, pertaining to the responsibilities of designers may take place.

One key ethical question we should be asking with respect to user measurement is not merely "Is it ethical to collect more information about a user?" (though of course in some situations that is the relevant question), but rather, "What information about the user are we not measuring, that we have a moral obligation to measure?"

What are the right things to measure? One is potential vulnerabilities on the part of users. This includes not only signals that a user might be part of some vulnerable *group* (e.g. children or the mentally disabled), but also signals that a user might have particularly vulnerable *mechanisms*. (For example, a user may be more susceptible to stimuli that draw them into addictive or akratic behavior.) If we deem it appropriate to regulate advertising to children, it is worth asking why we should not similarly regulate advertising that is targeted to “the child within us,” so to speak.

Another major area where measurement ought to be advanced is in the understanding of user intent. The way in which search queries function as signals of user intent, for instance, has played a major role in the success of search engine advertising. Broadly, signals of intent can be measured in forward-looking forms (e.g. explicitly expressed in search queries or inferred from user behavior) as well as backward-looking forms (e.g. measures of regret, such as web page “bounce rates”). However, the horizon of this measurement of intent should not stop at low-level *tasks*: it should include higher and longer-term user *goals* as well. The creators of technologies often justify their design decisions by saying they’re “giving users what they want.” However, this may not be the same as giving users “what they *want* to want.” To do that, they need to measure users’ higher goals.

Other things worth measuring include the negative effects technologies might have in users’ lives – for example, distraction or decreases in their overall well-being – as well as an overall view of the net benefit that the product is bringing to users’ lives (as with Couchsurfing.com’s “net orchestrated conviviality” metric).¹⁵ One way to begin doing this is by “measuring the mission” – beginning to operationalize in metrics the company’s mission statement or purpose for existing, which is something nearly every company has but which hardly any company actually measures their success toward. Finally, companies can measure the broader effects of their advertising efforts on users – not merely those effects that pertain to the advertiser’s persuasive goals.

Ultimately, none of these interventions – greater transparency of persuasive design goals, the development of new commitment devices, or advancements in measurement – is enough to create deep, lasting change in the absence of new mechanisms to make users' voices heard in the design process. If we construe the fundamental problem of the attentional economy in terms of attentional *labor* – that as users we're not getting sufficient value for our attentional labor, and the conditions of that labor are unacceptable – we could conceive of the necessary corrective as a sort of “labor union” for the workers of the attention economy, which is to say, all of us. Or, we might construe our attentional expenditure as the payment of an “attention tax,” in which case we currently find ourselves subject to attentional taxation without representation. But however we conceive the nature of the political challenge, its corrective must ultimately consist of user representation in the design process. Token inclusion is insufficient: users need to have a *real* say in the design, and *real* power to effect change. At present, users may have partial representation in design decisions by way of market or user experience research. However, the horizon of concern for such work typically terminates at the question of business value; it rarely raises substantive political or ethical considerations, and never functions as anything remotely like an externally transparent accountability mechanism. Of course, none of this should surprise us at all, because it's exactly what the system so far has been designed to do.

I'm often asked whether I'm optimistic or pessimistic about the potential for reform of the digital attention economy. My answer is that I'm neither. The question assumes the relevant task before us is one of prediction rather than action. But that perspective removes our agency; it's too passive.

Some might argue that aiming for reform of the attention economy in the way I've described here is too ambitious, too idealistic, too utopian. I don't think so – at least, it's no more ambitious, idealistic, or utopian than democracy itself. Finally, some might say “it's too

late" to do any or all of this. At that, I can only shake my head and laugh. Digital technology has only just gotten started. Consider that it took us 1.4 million years to put a handle on the stone hand axe. The web, by contrast, is fewer than 10,000 *days* old.

NOTES

- 1 Taylor, Charles (1992). *The Ethics of Authenticity*. Cambridge, MA: Harvard University Press.
- 2 Richards, Jef I. and Curran, Catharine M. (2002). Oracles on "Advertising": Searching for a Definition. *Journal of Advertising*, 31 (2), 63–77.
- 3 Paine, Thomas (1776). *Common Sense*.
- 4 MacManus, Theodore F. (1928). The Nadir of Nothingness. *Atlantic Monthly*, 19 (May), 594–608.
- 5 Rorty, James (1934). *Our Master's Voice: Advertising*. New York: John Day. Rorty's title refers to *His Master's Voice*, the famous painting of a terrier listening to his dead master's voice being replayed on a wind-up gramophone. The phrase later became the name of a British record company; today, both the phrase and the image persist in the name and logo of the entertainment retail company HMV.
- 6 Packard, Vance (1957). *The Hidden Persuaders*. London: Longmans, Green, p. 159.
- 7 Levy, Neil (2007). *Neuroethics*. Cambridge University Press.
- 8 Manjoo, Farhad and Roose, Kevin (2017). How to Fix Facebook? We Asked 9 Experts. *New York Times*, September 11. www.nytimes.com/2017/10/31/technology/how-to-fix-facebook-we-asked-9-experts.html
- 9 Schwartz, Jason (2017). Tagging Fake News on Facebook Doesn't Work, Study Says. *Politico*, September 11. www.politico.com/story/2017/09/11/facebook-fake-news-fact-checks-242567
- 10 Marcus Aurelius, *Meditations*, 8.14.
- 11 Popper, Karl (1970). The Moral Responsibility of the Scientist. *Induction, Physics and Ethics* (pp. 329–336). Dordrecht: Springer.
- 12 Rotblat, Joseph (1995). Nobel Lecture. Remember your Humanity. The Nobel Foundation. www.nobelprize.org/nobel_prizes/peace/laureates/1995/rotblat-lecture.html
- 13 Designer's Oath. <http://designersoath.com/index.html>
- 14 Hippocrates, *Of the Epidemics*, Book I, section 11.5.
- 15 Edelman, Joe. www.Couchsurfing.com