

Value Capture

C. Thi Nguyen

This is a pre-print. Please cite the final version,
Forthcoming in *Journal of Ethics and Social Philosophy*

Value capture occurs when an agent's values are rich and subtle; they enter a social environment that presents simplified — typically quantified — versions of those values; and those simplified articulations come to dominate their practical reasoning. Examples include becoming motivated by FitBit's step counts, Twitter Likes and Retweets, citation rates, ranked lists of best schools, and Grade Point Averages. We are vulnerable to value capture because of the competitive advantage that such crisp and clear expressions of value have in our private reasoning and our public justification. There is, however, a price. In value capture, we take a central component of our autonomy — our ongoing deliberation over the exact articulation of our values — and we outsource it. And the metrics to which we outsource usually engineered for the interests of some external force, like a large-scale institution's interest in cross-contextual comprehensibility and quick aggregability. That outsourcing cuts off one of the key benefits to personal deliberation. In value capture, we no longer adjust our values and their articulations in light of own rich experience of the world. Our values should often be carefully tailored to our particular selves or our small-scale communities, but in value capture, we buy our values off the rack. In some cases — like decreasing CO₂ emissions — the costs of non-tailored values are outweighed by the benefit of precise collective coordination. In other cases, like in our aesthetic lives, they are not. This suggests that we should want different values suited to different scales. We should want value federalism. Some values are perhaps best pursued at the largest-scale level, others at smaller scales. The problem occurs when we exhibit an excess preference for the largest-scale values — when we consistently let the universal metrics swamp our quieter interests.

Here is a story about how metrics can change people. A relative of mine had been planning a long European vacation with some old friends, John and Shelley. My relative had been looking forward to seeing the sights with her friends — touring the museums, seeing operas, having long dinners. But, she says, the entire vacation was dominated by John and Shelley's relationship with their FitBits. John and Shelley wouldn't go to the opera with her: not

enough steps. They'd cancel dinner dates because they hadn't met their daily step-goals yet. My guess is that that John and Shelley never consciously decided that step-counts were more important than, say, art or friendship. The FitBit just spoke more loudly in their internal deliberation, and there was no ArtBit or FriendBit to compete. The clarity of those metrics just swamped quieter considerations.

And, even if fitness was your main goal, the FitBit can exert a narrowing influence. Exercise can be valuable in all sorts of ways that aren't measured by a FitBit. A FitBit doesn't capture the ecstasy of complex skillful motion. It doesn't capture the camaraderie of team sports, the meditative calm of paddling a canoe across a quiet lake, or the aesthetic loveliness of a delicate rock climbing move. A FitBit measures exactly one thing: steps. That limitation arises from its particular institutional and technological embeddedness. FitBits are constrained by what mass-produced devices can easily measure and aggregate, given present-day technologies and institutional arrangements. We know how to make a watch that automatically measure steps, but not how to make a watch that automatically tracks your spiritual renewal.

Of course, you don't have to value what the FitBit measures. You could just use a FitBit as a source of data. But the FitBit tempts us to do more. The FitBit presents its output, not just as mere information, but as an evaluation: a score. And when you buy into the FitBit's preferred motivational scheme — when you adopt its scores as your values — you get all kinds of rewards. You gain the motivational benefits of having clear feedback about how well you're doing, of competing along a well-defined scale. All you have to do is give up on having fine and detailed control over your own values. Here's one way to put it: when you buy into a FitBit's preferred value system, you are *outsourcing the process of value deliberation*.

FitBit is just one example of a larger phenomenon which we can call *value capture*. Value capture happens when your environment presents you with simplified versions of your values and those simple versions come to dominate your practical reasoning. Value capture offers you a quick short-cut – an opportunity to take on *pre-fabricated* values. You don't have to go through the painful process of value deliberation if you can get your values off the shelf.

I want to focus on one particularly clear, and quite common, form of value capture: when an institution that presents you with some metric, and then you internalize that metric. You start exercising for your health, but you come to care about losing weight or optimizing your Body Mass Index. Or: you go on Twitter to connect to people and have fun, but come to care more about maximizing your Like, Retweet, and Follower counts. Or: you go into philosophy graduate school for a love of wisdom, but come out aimed at getting fancy grants, publications into highly ranked journals, and placement at a highly ranked institution. As anthropologist Sally Engle Merry puts it, the culture of indicators and metrics is “a form of governance that engages a person in governing himself or herself in terms of standards set by others” (Merry, 2016, 33). I will focus, for much of this paper, on such institutional value capture. Metrics are the starkest case of value capture, and we are fortunate to have a rich empirical literature studying the social effect of metrics. But metrics are just a starting point; there are many other forms of value capture worth investigating.

Many of us feel an intuitive horror when contemplating cases of institutional value capture. But it is rather difficult to say, in a principled way, exactly why value capture is so horrifying. For one thing, value capture is often consensual. People buy FitBits precisely because they know that those step counts will motivate them; they *want* to be captured, because the motivational bump seems worthwhile. Such gamified technologies are frequently sold as a

way to overcome weakness of the will, and seem to succeed at doing so. The point of a FitBit is to motivate you to walk more, and it does seem to work.

Why might this strike some of us as horrifying, rather than as simply a useful and empowering tool? I will suggest that there is a problem with the *nature of the values* on offer. The problem with internalizing institutional metrics isn't simply that we are getting our values from the outside. It is that such metrics are subject to the demand for a certain kind of stability and institutional usability. These institutional demands pushes our metrics away from the subtle, the dynamic, the sensitive — and towards what can easily be measured at scale, propagated across institutional units, and recorded in institutional memory. When we take on such metrics as our values — when we internalize them — we are imposing a narrowed filter on our values. We are letting the logic of institutions play a determining role in the articulation of our values.

Institutional value capture offers us a delightful reward. Once we have permitted ourselves to be value captured, our values become clear, coherent, and shared. Now we can be easily understood — unambiguously, almost effortlessly. But such clarity requires a degree of stabilization. Such clear, stabilized values arise from, and are deeply embedded in, external institutions and institutional processes. That stabilization has some benefits and some costs. Sometimes, those costs may be worth paying – but we should at first get clear about what, exactly, they are.

In value capture, we *outsource the process of value deliberation*. And, as with other forms of outsourcing, there is a trade-off. You get the outsourced objects quickly and easily, and they fit neatly into a larger network of other standardized and modular parts. Somebody else

has formulated our values for us, and done the work of embedding them in readymade systems of measurement and technologies of motivation. When we adopt those values, we gain access to readymade methods for justification. It's easy to justify yourself in the language of metrics, because metrics are easy to understand. They have, in fact, been engineered to be so. The cost of value capture is that we give up on the process of finely tuning our values to our own context: our personalities, our peculiar culture, our particular corner of the world. Outsourced values are not custom tailored. In value capture, you're getting your values off-the-rack.

A case study: the law school rankings

The social draw of quantification has been the subject of some extremely useful empirical recent studies, from anthropologists, historians, and sociologists. My favorite is *Engines of Anxiety* (2016), a study of the cultural effects of the *US News & World Report's* law school rankings, by sociologists Wendy Espeland and Michael Sauder.

Before the *USN&WR*, they say, there were no law school rankings. Students often picked law schools through a complex process of evaluation, deliberation, and self-reflection. They got to know a school by reading about its mission, by talking to people, or by visiting. Importantly, different law schools pursued different missions. Some were tuned to academic legal research, others to the corporate world. Some law schools were devoted to social activism — towards supporting the local community or serving underrepresented populations. The process of choosing a law school often triggered a certain degree of soul-searching in the students. The complex value plurality involved in the choice pushed students to reflect on

out what they wanted from their own legal education and legal career.

The rankings displaced all that. Espeland and Sauder studied online discussions between prospective law students. They found that, once the *USN&WR* started publishing its rankings, those rankings came to dominate the choice process for most students. And the same is true for non-students: the public perception of law schools immediately re-oriented itself along the *USN&WR*'s rankings. Espeland and Sauder say that the rankings drove value plurality out of the legal educational system. Many schools used to genuinely pursue their different missions. And many of those missions involved pursuing values that aren't tracked by the *USN&WR*'s ranking formula — like, say, supporting local underserved minority communities. But following such a distinctive mission invariably meant dropping spots in the rankings, which promptly resulted in precipitous drops in donations and student interest. Most schools, report Espeland and Sauder, have since abandoned their original missions, and re-oriented their admissions process and educational methodology towards performance in the *USN&WR*'s ranking calculations. And what matters the most to that ranking is the GPA and LSAT score of the incoming class, and the employment rate of the outgoing class (43).

In the case of the law schools themselves, the change in goals could be understood as a case of perverse incentives. Law school administrators were forced to align their efforts with the rankings, even if their own values were unchanged. But with prospective law students, the problem seems to run much deeper. The rankings seem to exert a magnetic pull over students' values. Some students, of course, were merely responding to incentives — since potential employers also care about law school's rankings. But a majority of students, say Espeland and Sauder, seemed to care directly about those rankings. Instead of exploring their own values and desires for their legal education, they seem to presume that the process

of going to law school should be oriented towards getting into the “best” law school, where “best” is determined strictly by the rankings. The existence of that clear, vivid, objective-seeming list offers an easy substitute for the process of personal value deliberation.

The effect on students I take to be a clear example of value capture. The fact that value capture exists I take to be an empirical matter — and its existence is well-documented.¹ My goal here is to think about the harms of value capture.

Value Capture

Value capture happens when:

1. An agent has values which are rich, subtle, or inchoate (or they are in the process of developing such values).
2. That agent is immersed in some larger context (often an institutional context) that presents an explicit expression of some value (which is typically simplified, standardized, and/or quantified).
3. This explicit expression of value, in unmodified form, comes to dominate the entity’s practical reasoning and deliberative process in the relevant domain.

If you’d like a portable version, try this: *value capture happens when a person or group*

¹ Beyond Espeland and Sauder, see Porter (1995), Scott (1998), and Merry (2016) for good entry points into the literature.

adopts an externally-sourced value as their own, without adapting it to their particular context.

Let's take a moment to get clearer on what, exactly, counts as value capture. First, notice that value capture includes both voluntary and non-voluntary adoptions of an external value. It certainly counts as value capture if, say, you were brainwashed and an external value was somehow injected into you, against your will. But it equally counts as value capture if you willingly and voluntarily adopted that external value – perhaps because it is easier, or helps you to fit more easily with in your profession, or because it lets you avoid the painful process of value deliberation. The target of my criticism here is not simply those cases of involuntary value transformation. I am interested what the problem with letting externally-sourced values dominate one's practical reasoning – even if that dominance was established knowingly and consensually.

Next, my definition of value capture is narrowly aimed at those cases where the entity uses the external expression of value precisely as given. It is aimed at those cases where we internalize and deploy an external value just as we found it, *without further adjustment* – without further contouring it, interpreting it, or fine-tuning it to ourselves. Value capture does not include cases where you get the seed of your values from the outside, and then start fiddling with them. If you get the starting seed for your values from your family, your culture, your religion, but then tweak them to fit your personality and place in the world – that's not value capture. Value capture is when an externally-sourced value, like a metric, comes to dominate your practical reasoning, in its given form – when your goal is simply to get to that higher ranking, those higher citation rates, those more Likes.

I've been speaking so far about the value capture of individuals by large-scale institutions. Such examples are vivid and familiar. But they can invite a simplistic reading of the problem: that real values are somehow original inventions of the individual, and that socially generated values are somehow fake. It's tempting to think that what's going on here is, say, a battle over individual authenticity, some conflict between the solitary free spirit and the forces of social conformity. But the problem is much more complex than that. For one thing, our values are often acquired, in their initial seed, from social sources — parents, teachers, friends, colleagues. For another, we often develop our values in community with others.

And crucially, value capture is a problem that can afflict groups too. A philosophy department can be captured by the larger university's focus on student evaluation scores. In my own experience, the clarity of an institutional metric can quickly come to dominate the attention of a deliberating group. Even when a group agrees that they care more about some inchoate value — like, say, fostering curiosity — the actual day-to-day decisions end up driven by whatever clear metrics happen to be on hand. Merry, in *The Seductions of Quantification* (2016), offers a good example. When the United Nations started publishing the Human Development Index — a quantified ranking of all countries in the world, in terms of how they supported quality of life. Merry says that the committee behind the HDI published it with a very clear and loud set of qualifications. They published with a lengthy report on the complexity and multidimensionality of “quality of life” measures, and clearly stated that the HDI ranking was simply a gross oversimplification. Unsurprisingly, says Merry, the fuller report was largely ignored. Once the HDI was published governments the world over became incredibly invested in advancing their ranking — even though the score wasn't attached to

any concrete real-world incentives or rewards. Here is a case where entire governance cultures have been value captured by an external metric. So: in my account of value capture, I intend “agent” to be in a broad sense, including individual persons and group agents.²

Next: 3.) specifies that the external expression of value “comes to dominate” the entity’s practical reasoning and deliberative process in the relevant domain. I mean the notion of “dominate” to be quite substantive here. Value capture occurs when an external value becomes the dominant source of reasons for action in a domain. It is not value capture if I adopt an external value in a controlled manner – as temporary instruments, accountable to my own richer values. I mean to exclude here, from the category of value capture, those cases where we use external values as proxies and heuristics under full reflective control – when we select, monitor, and adapt those heuristics in the light of our own richer values.

Suppose that I want to get healthier and more fit. By “healthy and fit”, I mean something complex and textured and difficult to express – something about feeling good in my body, being more capable of comfortably executing complex physical tasks, and getting rid of this feeling of awkward clumsy brokenness that too much laptop time has left me with. But such inchoate and airy expressions of value are pretty hard to use in the rush of daily life. Beings like us need heuristics – simple and clear rules of thumb to use in the day-to-day. And I can pick a heuristic, like increasing my step counts, as a quick-and-easy decision procedure to use in my daily life, as a way of pursuing that richer notion of health.

But such heuristics aren’t usually supposed to supplant our fuller values entirely. We

² I use “entity” rather than “agent” here because, while “agent” includes group agents, I think the category is not large enough. I suspect that some loosely organized communities can qualify as having values, but not have sufficient internal cohesiveness to count as a group agent. See discussion of shared values in Hedahl & Huebner (2018) and of loose community values in (Nguyen and Strohl, 2019).

are supposed to use them with the knowledge that they are mere proxies for our fuller values. They are supposed to serve our dominant values, which means they should be revisable and discardable under the light of our fuller values. What I really want is health in this richer sense, but I also I know that I need an easier target to aim at on daily basis in order to get myself motivated. So I start using a FitBit, and just aim at getting step counts. But after a few months, I step back and reflect on my time with the FitBit. Has pursuing step-counts made me happier? Is my body performing better? Do I feel less broken and awkward? Perhaps the answer is affirmative and I keep going with the FitBit; perhaps the answer is negative and I abandon it and try some other proxy goal. Perhaps the answer is a qualified yes and I modify my approach, adding a few more goals to the mix, beyond just maximizing my steps. This *controlled use of a proxy value* is not a case of value capture, since the externally-sourced value doesn't dominate my deliberative procedure. The dominant value is not the FitBit's step measures but something else, and this can be seen by the fact that I do sometimes adopt a reflective stance where I decide whether adopting the FitBit's goals are serving my real values, and decide whether to continue or discontinue my use of that simple proxy in the day-to-day. The FitBit is not in charge.

Similarly, it isn't value capture when I am merely *taking the metrics into account* in the pursuit of my own rich and textured goal. Suppose I want to be a legal activist working for immigration reform. I know that going to a high-ranked law-school will be important for getting the influence I need for this kind of work. In that case, I'll pay attention to the law school rankings – but they don't dominate my practical reasoning. I may attach an instrumental value to going to a highly ranked law school, but I can also trade off that ranking in light of my real values. I won't, for example, just go to the highest ranked school possible. I'll use the

rankings for my purposes – trying to find a compromise between a well-ranked school that will get me the power I need and a school that will help me learn to do the activist work I want to do. That kind of instrumental awareness of ranking systems is a long way off from a real case of value capture – where somebody’s primary goal was, say, simply to go to the highest-ranked law school.

Next: value capture can happen at different points in an agent’s life arc with values. Sometimes, an agent has already established their values, and then they come to be *replaced* by some external metric. Other times, the agent doesn’t yet have their own articulated values; they are in the process of figuring them out. But the existence of a pre-fabricated value offers them a *shortcut* in the process of value deliberation. They can simply adopt a ready-made value instead of going through the slow and oftentimes painful process of figuring out and adjusting their values to their own personality and circumstances. The definition of value capture is intended to include both replacement and shortcut cases.

Finally: value capture can happen at some different loci One kind of value capture involves the *wholesale* capture of the entire value — such as when you got into this career for joy, but come to care only about the money. In wholesale value capture cases, the agent changes how they think of their values in their space; they come to describe their values differently and report them differently. But just as common as these wholesale cases, I suspect, are cases of what we might call *application capture*. In such cases, an external expression of a value doesn’t replace how we conceive of our original value – in how we would think about and report our values in the abstract. But the external value dominates how we act by

setting the practical criteria in day-to-day applications of our values in particular decisions and evaluations. Say that I, an academic, care about the pursuit of truth, wisdom, and understanding. Across my career, if asked, I would describe my core values using those same terms. But suppose that, of the course of my professionalization, the way I apply those terms changes. Now, whenever I try to evaluate the success of my articles, I turn to certain metrics, like the citation rate or the status of the publication venue on some ranked list. And when I evaluate my overall success as an academic, I turn to metrics like my total citation rate or the status of my institution on some ranked list. In that case, it is those institutional metrics, and not the vaguer values I report upon reflection, which effectively dominate my actual actions and self-evaluation. Here, the metric gains dominance by capturing, not the general terms in which I articulate my values, but the more specific application criteria I use when the values hit the ground. The metric fills out the process by which I determine whether I have fulfilled my core values. And suppose that I guide my actions based on those evaluations: I start writing papers that are more like the ones that have succeeded, in these terms, and start taking actions that might advance my general success, in these terms. Then those external criteria have come to effectively dominate my practical decisions. There is a crucial difference between controlled use of a proxy, as I described it earlier, and application capture. When we *use* a metric as a mere proxy, our richer values are in charge. We will regularly reflect on the proxy from the perspective of our fuller values, and modify, discard, or adapt that proxy. In the application capture cases, we let the proxy take charge. It functions as the effective functional translator, connecting our abstract expressions of value into specific cases of evaluation – controlling how we apply our values to the world. (Though, at least from my own observations, many cases of value capture start as innocuous-seeming uses of a proxy. I've

heard from many people say that they put on a FitBit in order to pursue some other goal, like health or happiness, but then years later they found that they had forgotten about that larger goal -- that doing well in the FitBit's terms had come to occlude all else.) From here on out, I will speak of "values" being captured, for brevity's sake — but I mean to indicate both whole-sale value capture and application value capture.

To sum up: value capture does not include every interaction with rankings or metrics. It doesn't include the controlled use of proxies and heuristics, nor the informational use of metrics. Value capture occurs when an externally-sourced value plays the dominant role in practical reason – when it gets put in charge, for some domain. This looks like: people who pursue step-counts even when it hurts their knees and exhausts their spirit; academics that pursue publications in the highest ranked journals even when their work feels boring and meaningless; universities that pursue high rankings in the *US News & World Report* over richer understandings of education; newspapers that pursue clicks and pageviews over their own sense of newsworthiness and social importance. And, as I've noted: the empirical work indicates that this sort of robust value capture is actually quite common.

Value capture is distinctive because we don't change or adapt the particular externally-sourced specification of a value to our particular context. Compare this with other, more open-ended and dynamic relationships we might have to externally-sourced values. We often get our first grip on a new pursuit – and its value – with another's help. A friend shows me the wonders of horse-riding, the beauty of jazz, the depth of haiku. They talk about what they find meaningful and rich in the activity; they guide me into the actions and attentions that will help me get onto its distinctive value. As Tal Brewer (2009) puts it, the values of activities are often obscure to the outsider or novice; it takes a long process of immersion in

the activity to get onto its true value. And we often need help to find our way in.³ The friend who taught me to see jazz talked me into it — into the particular thrill of seeing a live improvisation. But I suspect she would have been very disappointed if, ten years down the line, the value I found in jazz precisely mirrored her own. Like any good art friend, she hoped that I would eventually fly on my own wings, and sharpen the details of my love of jazz in my own way.

And I did – I used her guidance to find my way in, and then slowly began to develop my own relationship with jazz, finding out what thrilled and moved me in the music. This isn't a case of value capture; I have used external guidance to get my first grip on the terrain of value in an activity, but then significantly tailored my sense of value. Value capture are the cases where I internalize, wholesale, an externally sourced value and permit it to dominate my reasoning in its unchanged form. This is where the “outsourcing” metaphor is particularly useful. The harms of outsourcing do not depend on any involuntariness. I can wholeheartedly consent to outsourcing. The harms come from the particular content and nature of outsourced objects – of their inflexibility and pre-fabrication. If you want a slogan: our values should be tailored to our particular selves and our particular context – but in value capture, we buy our values off-the-rack.

The problem of value capture

What, then, is the cost of outsourcing one's values? First, to be clear: I am not trying to

³ See Nguyen (2021a) for a discussion of how trust in others is often required to provide the motivation for attending to difficult or obscure art forms.

argue that value capture is always wrong. Value capture, as with any other form of outsourcing, involves a trade-off between efficiency and fine-tuning. I think that we are often clear on the benefits of that trade-off, but fail to plainly see the costs. My goal is to articulate more precisely the costs, so we can get clearer about what we're paying, and whether it's worth it.

But it turns out to be rather hard to articulate the problem with value capture. First, what is wrong with getting our values from external sources? It seems utterly naive to think that our values need to spring fully formed some magical inner place, wholly devoid of social origin. We are deeply social beings, and we often seem to get our values from our culture, our community, our social context. Second, how could value capture undermine autonomy? Many cases of value capture are entirely voluntary and consensual. They know that FitBit motivates because it presents information in public and shared terms – though the full implications of that publicity may not be entirely obvious. Consensual value capture can seem like an aid to autonomy. People often seek out such gamifications in order to overcome weakness of the will. People buy FitBits or use DuoLingo precisely because the gamified structure — in which they are awarded points and levels for progress — makes them more able to get certain things done, like start exercising or learn a language. They are *hoping* for value capture, and they are choosing the effect with some understanding of the basic mechanism. As Jane McGonigal (2011), puts it, gamification is a force for good, because it can turn monotonous tasks fun. If value capture can help us overcome weakness of the will, then it helps increase our autonomy and agency. So what's the harm?

There are at least three ways to think about the potential harm of value capture. First, it might be that *autonomous* participation in the formulation of our values seems good in and of itself – and not just mere one-off consent to a big package, but a fine-grained and ongoing

autonomous control of the details of our values. If that were true, then value capture would undermine our autonomous control over our values.⁴ Second, institutional values are subject to demands for *hyper-explicitness*, and hyper-explicit values seem unlikely to adequately capture the full richness and subtlety of human values.⁵ Third, the kinds of external values we encounter are typically formulated according to the interests and perspectives of large-scale institutions. They are, we might say, *standardized* values. Such values seem unlikely to fit the varying and peculiar interests and situations of particular people and smaller-scale groups.

A full account of value capture would need to address at least these three approaches, and I hope one day to provide such a fuller account. Here, I can only take a first step. I will concentrate on the issue of *standardized values* — in part, because I think it mostly sharply highlights the unique problems of value capture.

Here is the worry in a nutshell. Value self-determination is important for all sorts of reasons. Here's one: value self-determination yields values that are finely tuned to our particular context. By substantively participating in the detailed process of formulating our values for ourselves, we can get values that nicely fit our particular circumstances — our individual psychologies and phenomenology, our group culture, our local context. Value capture intrudes on that process of value self-determination, *substituting pre-fabricated and standardized values for finely tailored ones*. Note that this is not an argument that autonomous

⁴ Of the three options outlined here, this is the one I am most undecided about. While intuitively appealing, developing such an account depends on walking the tightrope between specifying a substantive condition of autonomy, while keeping a grip on the social sourcing of many of our values, even in the most autonomous cases.

⁵ I am exploring this possibility in other work. See Nguyen (forthcoming) for a discussion of the possibility that hyper-explicit values represent a bad epistemic attitude towards the world of value — that they discourage exploration of the space of value, by making it easy to dismiss new candidates for value.

value formulation is a good in and of itself, but rather an argument that substantively participating in the process of shaping one's values is instrumentally good in that it yields better, more finely tailored values.

Here is an example from my own life. For the first two decades of life, I avoided most physical activity. I had an incredibly simplistic conception of the value of exercise. I thought that exercise was basically pounding out some miles on a treadmill to burn some calories. Eventually, I came to see the vast and varied joys of athleticism. But in order to get there, I took a long meandering journey through many different sports, each of which paid off in profoundly different ways. Long-distance running turns out to be zen-like and calming. Trail-running requires more attention, but offers this thrilling sense of reactive flow to the difficulty of the trail. Deadlifting is brutal and intense, a pure shot of grueling focus. And rock-climbing turns out to be a fascinating fusion of bodily aesthetics and puzzle-solving, where you solve thorny movement puzzles through elegant motion.⁶ And even inside one of these activities, there isn't some singular value on offer. Rock climbing can be pursued in radically different ways, each of which offers very different rewards. You can seek out thrills and risk; you can do easy climbs in rapturously beautiful terrain; you can focus on finding climbs with graceful movement; or you can go for gruelingly athletic climbs on a cave roof just eight feet off the ground. Each of these different ways of valuing rock climbing suggests a different way of approaching it, which in turn yields richly different textures of activity.⁷ This is a process of *exploration*, where you try out things out, figuring out how they fit with you, and changing

⁶ For more on the aesthetic qualities of movement in rock climbing and other games, see Nguyen (2020a, 2020b).

⁷ This description has been deeply influenced by Tal Brewer's (2009) account of how the formulation of the value of an activity, and the way we do an activity, form a feedback loop, as we explore and refine our understanding of the activity. Agnes Callard (2018) has also written on such proleptic ends, though her account adds a requirement that the process is triggered by a desire to become the kind of person who so values.

around your approach in response, seeing how it goes in an ongoing loop of feedback and adjustment.

The worry is, then that when you are value-captured by a FitBit, you don't go through that process of exploration and fine-tuning. My claim here is not that one puts on a FitBit and is automatically value-captured. One could simply use a FitBit as a data-gathering system to pursue one's own values. But FitBit does presents its step-counts as a *score*. It is a gamified system, which openly employs design features from games.⁸ A FitBit doesn't force value capture, but it certainly invites it.

We might even call this an extended value system. Some philosophers have been very excited to claim that our minds are extended beyond our bodies — that our minds can include various technologies as parts of their internal functioning.⁹ Most of the discussion of extended mind has focused on adopting, as part of our extended mind, various value-neutral cognitive resources — like using a notebook or Google Docs as an extended memory. Some of the discussion has gone so far as to suggest that we can extend our mind to use various technologies as part of our emotion-regulation system, such as Joel Krueger's (2019) suggestion that we use our portable music-devices for mood-regulation. My suggestion is one further step: in some cases, a *standard for evaluation* is embedded in a technology, as in FitBit's step-counts or Twitter's Likes. When we integrate that technology into our cognition, our extended mind now includes a *value system* which was created externally, and which is sustained through external technologies.

⁸ See Nguyen (2021b) for a discussion of how certain user interfaces can present metrics as scores, and thus as forms of evaluation.

⁹ Alternately, if we want to avoid the endless tussles about what exactly is the line between mind and not-mind, we can use Kim Sterelney's (2010) locution: that various technologies are scaffolds for agency. Either locution takes us to the same worry.

Of course, one might respond, we get our value systems from external sources all the time — from our parents, our community, our culture. But, in the unproblematic cases, we can use external values as a starting seed, which we can adapt and tailor to ourselves. The worry is that we simply plug in these external values and use them as-is. In particular, some external values can *resist further tailoring*. This is especially likely when adopting a particular pre-fabricated value is appealing precisely *for the standardization*. Then we will be quite tempted to leave them as they are, or else lose out on the promised efficiency. Value standardization is like any other kind of standardization: we gain in efficiency, but in exchange for giving up localized tailoring (Bowker and Star, 2000).

To really understand the problem here, we need have a good grip why we should want to tailor our values to fit. Elijah Millgram offers a useful account of how we adapt and improve our values in *Practical Induction* (1999).¹⁰ He's not talking about abstract, generic renderings of value, like, say, "happiness" or "flourishing". He is interested in the specific, grounded articulation of our values and goals by which we conduct our day-to-day lives: such as a runner's pursuit of a better marathon time, or a filmmaker's interest in playfully subverting genre conventions, or a philosopher's interest in writing deep, rigorously argued papers. Crucially, says Millgram, we don't derive these specific articulations of value by deriving them from some abstract specification of the good. Rather, we acquire our particular values and goals via a process of practical induction. *We try on* particular goals and values for awhile. We might enter a profession and try on the goals associated with that profession: a literary fiction writer might start caring about achieving realism of character and setting; a Montessori

¹⁰ Millgram has since developed some of the ideas in *Practical Induction* further, most notably in argument that boredom and disengagement is a signal that one's values and chosen roles are a bad fit (Millgram 2004). My treatment here relies on both of his discussions.

teacher about fostering autonomy in very young children. And then the person gets *feedback* from the experience of living life under that particular value system. They might get positive feedback, like feeling engaged, happy, or interested; they find themselves savoring the details of their life. Or they might get negative feedback: they feel bored, listless, disengaged. This is feedback about the *fit* between the values we have adopted and the particular circumstances of our lives: our personality, our culture, our place in society. To flourish, we need to be sensitive to that feedback, and use it in fine-tuning our values to fit.

Millgram's own discussion focuses on large-scale value shifts which accompany things like, say, having a mid-life crisis and changing career. But his argument leaves room for smaller-scale adjustments in the articulations of our value. Say I start rock climbing, and take up the most obvious standard of success in that hobby. There is a generally agreed-upon difficulty scale for climbs; most new climbers just start by trying to advance up that scale. Some climbers flourish under that goal; others do not. When I focused on advancing on the difficulty scale, I found myself miserable, tormented by my sense of inadequacy and my inability to progress. My climbing days filled with exhaustion and dread. So I began to change my sense of my goal in my climbing, started pursuing a slightly more personal vision. I started looking for the most elegant, interesting climbs, and my goal became to climb them as delicately and with as most control as I could. And under that goal, I flourished: I became more constantly sensitized to the details of my movement and more at peace with simply enjoying a bit of lovely climbing; rock climbing trips now left me feeling rested and happy.

So here is a first pass at the *tailoring argument* for the harm of institutional value capture. If Millgram is right, then we will flourish when we have the capacity to adjust and tailor our values in light of our rich experience of the world living under them. When we tailor our

values to ourselves in light of those rich experiences, then our values will be better fit to promote our flourishing, as the very specific people we are, in our very specific circumstances.

Perhaps you do not like the references to some ill-defined sense of “human flourishing” or “well-being”. We can put the same thought in less mysterious terms. When we adjust our values in light of our rich emotional experience of the world, then those adjusted values will be better suited to support a more emotionally positive life. If we adjust our values, taking interest and engagement as a positive sign, and boredom and ennui as a negative sign, then our values are more likely to give us a rich, interested, and engaged life, rather than a bored and listless one. But in institutional value capture, we do not adjust our values in light of our particular experiences. We take values as provide by some large-scale institution, and live under them as given. Those values will have been formulated to take deeply into account various institutional interests: like the ability to be counted in a reliable way across a large institution, and the ability to be readily aggregated in an institutional bureaucracy. They will not have been formulated in light of the rich feedback of how our particular lives have gone when we live under these values. In value capture, we adopt values that values have been formulated in way that is *insensitive to* and therefore *less able to support* our rich, subtle, and personal emotional experiences.

The first pass emphasized the problems of value capture for the individual. This reflects Millgram’s own version of the argument, which emphasized individual values and individual phenomenology. But we can also easily extend the argument to encompass group value capture. Groups, too, have particular articulations of their values. This can look like, say, the community of analytic philosopher’s value in rigor or creative writing community’s value in

personal expression.¹¹ The contemporary community of improv comedians, for example, have come to put a strong value in collaboration via automatic agreement. The core rule is “Yes, And...”: you always accept other people’s suggestions and build on them. This expresses a value, we might say, of *radical acceptance* – of never refusing ideas, and always integrating every proposal and building upon it. This value works extremely well in the context of improv comedy. And the precise articulation of this value has obviously evolved over the years, through trial and error in countless acts of improv comedy. But its success is context dependent. (Imagine trying to center such a value in analytic philosophy. Analytic philosophy, one might think, is a century-long social experiment in the value of harsh criticism and the radical refusal to accept anything.)

So long as there are accessible signs of a groups’ flourishing, or of a community’s well-being, then Millgram’s account of practical induction should also apply to the development of group values. Such group-level value tailoring is unlikely to center internal emotional phenomenology as strongly. But groups can tailor their values in response to their particular nature and context. Jane Jacobs (1961) offers a particularly vivid example of how we might tailor a specific value to a specific context. Dwellers in dense urban environments, she says, have learned to cherish *privacy* in a way that suburban and rural people do not. So much of one’s life is conducted in dense public environments, that city-dwellers have developed a profound devotion to maintaining privacy: of not making unnecessary eye contact, of not intruding into nearby conversations. Valuing a certain kind of eager “friendliness” – easy eye-contact, being willing to start conversations with anybody at anytime, makes perfect sense

¹¹ For an excellent discussion of the scientific community’s interest in enlarging the collective data supply, see Strevens (2020).

in lower-population density areas, without that constant press of humanity. But in a dense city, without that collective devotion to the practice of privacy, city-dwellers would be utterly overwhelmed by constant social interactions and demands.

To sum up: it is good for agents — individual and group — to tailor their values to their particular context. Those values will be better suited to support the well-being and flourishing of individuals, groups, and communities, by being adapted, in their formulation, to the particular nature of the agents, and to their particular context. Value capture interferes with that tailoring. It does so even when the value capture is the result of a fully informed and consensual process, since the problem lies in the *content* of the values, and not in the bare fact of their voluntary adoption.

Here's another way to put it: value capture, even when consensual, involves a low degree of granular control over the details of the contents of one's value. It puts you in the same relationship with your values as you have with, say, your iPhone's End User License Agreement. When you click to sign a EULA, you did, technically, consent, and you are, technically, responsible. But you only have one binary choice: accept the whole package or not. When we permit ourselves to be value-captured by institutional values, we have the same low granularity of control over our values: we either accept the whole package, or not. You can't get control over how your FitBit counts steps, or how the edifice of higher education counts citation rates and impact factors.

This low granularity arises directly from the core functioning of large-scale collective values. To better understand why, let's look at the processes that drive the creation of institutional metrics.

Metrics and the standardization of value

Let's focus on value capture by institutional metrics. I think this is the starkest case of value capture, and a good starting point for thinking about other forms of value capture. My goal, in this section, is to make clear why institutional metrics resist value tailoring as part of their essential functioning.

One response to the case studies I've offered so far — the *USN&WR* case, and the UN's human rights metrics case — is that the particular metrics are bad. Perhaps there's nothing wrong with value capture per se; it's just that we need to pick good metrics. But I'll suggest that we are unlikely to find any institutional metrics that are good to take on as individual, or small-group values. Metrics are formulated to serve certain key institutional interests – to work at large scale – and they need to be relatively inflexible to play their role. Institutions want metrics that are narrowly specified, standardized, and inflexible. Precisely what makes a metric good in the institutional context will make it problematic to internalize as a value for individuals and small-scale communities.

Here, we can turn to a rich and useful empirical literature on the place of quantification and standardization in bureaucracy and political life. Here, we are the beneficiaries of decades of empirical study of quantification culture, performed across a number of fields, including history, sociology, anthropology, and communications.¹² What follows may sound familiar to some philosophical ears; the empiricists I will be discussing are often working in

¹² The study of quantification culture is often associated with the interdisciplinary field called Science and Technology Studies.

a Foucaultian mode. The field, in particular, has been highly influenced by philosophical figures such as Ian Hacking, Bruno Latour, and Martha Nussbaum.

A foundational work here is Theodore Porter's (1995) history of quantification culture, *Trust in Numbers*.¹³ Porter is particularly interested in how quantified forms of justifications, like the cost-benefit analysis, came to dominate politics and management. He isn't arguing that quantification is always bad. Rather, his goal is to get clear on the relative advantages and disadvantages of qualitative and quantitative ways of knowing. Says Porter: Qualitative ways of knowing are nuanced and context-sensitive. But qualitative information is difficult to manage en masse and difficult to transfer across contexts. Qualitative evaluations usually require significant shared background knowledge to adequately interpret. When we transform information from a qualitative to a quantitative format, we strip off much of the nuance, texture, and context-sensitivity. By doing so, we create a *portable* package of information, which can be easily sent across contexts and understood by people with little shared background.¹⁴ Quantified evaluations can be easily transmitted between people with little shared background, precisely because they have been stripped of context-dependent features. And quantification isolates the more invariant parts of that information, so that the results can be readily aggregated. For this reason, quantitative methods are preferred by large-scale institutions, which must pass information across many levels of hierarchy — between distant administrators with low shared context (Porter, 1995, 3-86). In other words,

¹³ Though Porter is a historian, he was significantly influenced by Ian Hacking's work in the philosophy of science on the formation of categories and measures.

¹⁴ This notion of "portability" as the center of quantified information is alive in more contemporary work in this space. In Sabina Leonelli's (2015, 2016) crucial work on the philosophy of data, she defines "data" as information that has been prepared to travel to new and unexpected contexts.

quantifications are preferred in large-scale institutions precisely because of their narrowness and their context-invariant stability.

And quantitative evaluations themselves vary according to their nuance and context-sensitivity. Once, land in England was measured in *hides*. A hide is the amount of land required to support the average family. The hide is a measure which highlights a highly relevant functional quality. The acre is a measure of land size, rather than land function. Similarly, says Porter, older Polish land measures varied by soil quality, so a given unit of land would approximately represent a similar productive value (24). When a ruler attempts a fair distribution, the measure they use will determine which quality is evenly distributed — in this case, land size versus land functionality. Think about the difference between, say, a king's giving each of his soldiers a hide of land, versus his giving them each ten acres of land. One might think that the hide is a superior measure of functional worth, and so a vastly preferable measure for providing fair compensation. But hides are highly variable in size, and determining what counts as a hide requires the application of detailed local knowledge. A hide in a fertile river valley is smaller than a hide in a desert. Hides also vary depending on local weather patterns, game animal migration patterns and more. The hide is a measure that can really only be effectively managed at the administrative periphery — by locals, who know their environment and its inhabitants' typical needs and usage patterns. The hide is impossible to administrate from any sort of distant bureaucracy. So, says Porter, when we shift from small, local, distributed governance to large-scale centralized governance, we inevitably shift from informationally rich — but difficult-to-manage — measures like the hide, to more standardized, but informationally impoverished, measures like the acre.

James Scott calls this the state's view of the world. By 'states', Scott means any large-scale

institution, including governments, corporations, and the emergent networked institution of globalized capitalism. States, says Scott, can only manage what they can see, and they can only see that information which has been rendered into a form which can be processed bureaucratically — information that has been standardized and quantified. States can only see those parts of the world which has been rendered *legible* to them (Scott, 1998, 11-83).

Student grades provide a familiar example. In the modern educational environment, student grades are almost always quantified. But there are other modes of educational assessment. Imagine an educational environment where we only offered qualitative evaluation of their students' work, like written feedback describing its the good qualities and its problems. Such evaluations can easily pivot to address different dimensions – like the writing clarity, the originality, the argumentative clarity – without any demand to compress that all down to a single dimension of evaluation. Such evaluations can also be tailored to each student's own particular goals. I might give very different suggestions to a nursing student interested in the practical implications for their work than I would to, say, a future lawyer or future journalist. If our goal is simply to educate the student, we don't necessarily need to provide an overall rating of all our students on some single common scale.

But in our actual world, we must offer a quantified measure of each student's success — a measure which permits us to instantly compare any student with any other: their grade. This quantified ranking of students is extremely useful to administrators. All of a student's efforts in a class can be expressed in a single number. This also enables a further aggregation: all their class grades can be averaged to generate a single number, which represents their entire educational career — a Grade Point Average. And the existence of GPAs is enormously useful for the project of administrating a large-scale educational bureaucracy. They enable

all kinds of fast, easy, and objective-seeming manipulations. An admissions officer can arrange the data from every single student applications into a spreadsheet and quickly sort them by GPA. They can create an automatic cut-off point below which student applications are automatically discarded. Sets of student GPAs can be aggregated in order to yield a single number that can be used as a metric of performance for a particular teacher or a whole school district.

In their study of the history of American grading, Jack Schneider and Ethan Hutt argue that standardized grading schemes were implemented to make grades more legible and usable to administrators and employers (2014, 203). Before grading, there was no communicative “short-hand”. Evaluations required intimate communication between teacher and student. Early systems of grading were “low-stakes” affairs; they were set up differently in different schools and built to encourage student learning. But the modern system of grading serves, not a pedagogical purpose, but an organizational purpose. It enables students to easily transfer between different institutions. Perhaps most importantly it standardizes a product for future consumption on a market. Standardized grades make possible standardized educational certificates, which are extremely useful for potential employers. It was administrators and employers who “placed a premium on readily interpretable and necessarily abstract grading systems” (217-8). Qualitative evaluations of student might be nuanced and context-sensitive — but they are illegible to the large-scale administrative institution.

Finally, these various procedures — data-collection, transformation into standardized inputs, and aggregation — need to be codified into a set of policies that can reliably executed by very different people. Large-scale institutions need to train up people from different backgrounds to perform the same sorts of tasks. And their performance needs to be assessable

and auditable by others — where those auditors also come from different contexts, and their audit procedures themselves are subject to the same demands of explicability and transmissibility (du Gay, 2000). These various procedures need to be *standardized*. That means that the inputs and processing rules of these procedures need to be regulated across many contexts (Bowker and Star 2000, 13-16).

We can draw, from this mess of observations, some underlying themes. Institutions share a basic functional interest, inherent to the functioning of large-scale administrative systems. They need to manage information across a vast domain. This need arises intrinsically from the need for an institution to function as a coherent whole. Notice here that I am not presuming that the institution has some interest in controlling or manipulating individuals. Even the most well-intentioned of organizations — like, say, a charitable non-profit — has this same functional interest in information management. The interest arises from the basic conditions of coherent group agency, as instantiated in a policy-based, centralized bureaucracy.

This functional interest is served by two standard mechanisms — quantification and standardization. Institutions need to render the world into a format legible to large-scale institutional information-processing procedures. So institutions need information in quantified and standardized format. Because of their institutional function, these mechanisms — quantification and standardization — tend to share some specific features, which make them problematic to internalize as personal values. First, quantified metrics are *narrowed* by design. Only certain things count. Institutional measures need to be usable across different contexts. This requires that the measures leave aside highly context-dependent forms of understanding and focus, for their inputs, on context-invariant qualities. As Scott says, the narrowness of the metric creates a narrowness of institutional vision. Institutions can only

see, process, and act on parts of the world that are counted by their metrics. Anything that doesn't impinge on those metrics is invisible at an institutional level.

In value capture, we internalize those narrowed metrics, thus narrowing our values. And, insofar as our values drive our attention, then the value captured will be subject to an analogous effect to Scott's narrowed institutional vision. It's not that we literally don't see things that fall outside our narrowed values, but we won't devote much energy to them, or dismiss them as unimportant. Think here of the businessperson who thinks that only money matters, and who immediately dismisses from mind any unprofitable ventures — like art or philosophy.¹⁵

Next, such institutional metrics typically present values in highly explicated, *finished* form. They resist re-interpretation. Pre-institutionalized values are often expressed in an open-ended manner. A concept like "health" or "fitness" or "a good education" admits of different interpretations. Different people may work out their own interpretation of what counts as a good education — and so evaluate their understanding of the term. You want to know more useful things, I want to indulge my sense of curiosity — both are viable understandings of what one might want out of an education. But step-counts and law school rankings do not admit of such variability. The method of assessment is rigid. Says Porter, the process of quantification is useful to large-scale institutions, in significant part, precisely because it reformulates information so as to remove the need for interpretation (Porter, 1995, 21-29). Standardization is required for informational portability — and standardization requires rigidity.

¹⁵ I further develop this line of thinking in Nguyen (forthcoming), which explores the possibility that overly explicit articulations of value can narrow our attention and exploration.

And those off-the-rack values usually come *embedded* in institutional infrastructures, institutional language, and mass technologies, so as to resist further tailoring. We don't have the power to fine-tune the innards of such institutional values. We can't tinker with the way Twitter counts Likes, nor adjust with the *USN&WR*'s ranking algorithm. They are hard-wired into external systems. This rigidity and uncustomizability of measures and metrics is no accident. It is essential to their institutional function. Standardization enables easy communication and ready aggregation — but, to do so it must resist individual customization.

To summarize: institutional metrics are designed according to *alien interests*.¹⁶ They have, in fact, been fine-tuned and adjusted — but to satisfy interests that are not our own. I am not presupposing, here, that institutions must have malevolent intent, like an interest in domination, control, or power. To put Scott's discussion into philosophers' terms: all we need to attribute to an institution is a basic interest in *agency at scale* — an interest in gathering information about the world, managing that information, and using it to inform actions. But the scaled-up nature of bureaucratic institutions imposes certain distinctive requirements on that information-gathering process. Institutional metrics are typically formulated to fit the demands of scaled-up informational agency: for easy recording in institutional memory, for transmission across bureaucratic layers, and for manipulability by institutional methods. When we internalize institutional values, we are letting such interests play a powerful role in the formulation of our own values. Value capture gets us to take an institution's eye view on ourselves — to evaluate ourselves and our activities in institutional terms.

¹⁶ (Owen and Cribb, 2019) make a similar point in their analysis of FitBit technologies (32-35). They distinguish, however, between procedural autonomy — which involves internal deliberative processes — and substantive autonomy, which involves one's ability to actually act on and bring to fruition one's decisions. They say that, by and large, self-tracking technologies like FitBit may aid procedural autonomy, but cannot aid substantive autonomy, since such technologies can't fix large-scale social inequities. My argument is that such technologies also significantly undermine procedural autonomy.

We have much to gain by fine-tuning our values, fitting them with our psychology and world. Institutional metrics are tuned, not to an individual's rich and particular experience of the world, or a small community's particular context, but to the needs of information processing at a mass scale. In value capture by institutional metrics, our values become rigidly tied to an external expression. That rigidity arises, in significant part, from the institutions's interest in large-scale informational management. Metrics, by design, resist attempts at digestion and customization by the agent – and they usually come embedded in large-scale institutional infrastructures which make them even more inflexible. Their alienness resists adaptation.¹⁷

The seductiveness of metrics

Of course, we often have to use such metrics when we work within, and next to, institutions. But we could use them while also keeping them at emotional arm's length. We could employ them in our reports and our requests for funding, but only as the trade language of bureaucracies. Why might we ever take the further step and internalize them? The answer comes in several stages.

First, quantifications, in and of themselves, are seductive in their clarity and crispness.

¹⁷ One might ask what relationship this view has with various forms of alienation critique. Though my analysis here is obviously similar, in spirit, to the general themes of alienation critique, I avoid use of the term “alienation” because my analysis here differs, in key respects and in many details, traditional alienation critiques. As Rahel Jaeggi says, many forms of alienation critique involve views where the alienated agent is divided against themselves, as unable to identify with their work, as diffident and depressed (Jaeggi, 2014). But the value captured agent can be wholehearted (think of the capitalist all-in for money), fully identified with their work, energized and motivated. They are not divided against themselves; rather, they are simplified, where that simplification has been guided along institutional lines. Notice, furthermore, the difference between my analysis and the traditional Marxist alienation critique. It is possible to be value captured by a fully socialist bureaucracy. Here, I am aligned more with Scott's particular version of neo-Foucaultian critique, than with Marx. For Scott, both globalized capitalism and centralized communism share an interest in rendering the world legible into the terms which they can process and act upon.

Many people seem to trust quantified data simply because it is quantified. And we should certainly trust data when it has been generated using reliable methods. However, the mere quantified format itself often seems to generate trust, regardless of quality of the underlying methodology. But obviously, mere presentation in a quantified format does not offer any guarantee of reliability. So, insofar as we trust from the bare fact of quantified presentation, then that trust is unwarranted. And Porter, Merry, and Espeland and Sauder provide evidence aplenty that bare fact of quantification actually does, in fact, generate such unwarranted credibility. To put in the contemporary parlance, the excessive credibility given to quantified data counts as a form of epistemic injustice or epistemic oppression (Fricker, 2007; Dotson, 2014). It harms those who are unwilling or unable to present their information in such quantified form, preventing them from being appropriately recognized as sources of information. And insofar as quantified data tends to emerge from certain sorts of institutions, then those institutions themselves are the beneficiaries of epistemic injustice.

Why might the mere presentation of information in quantified form invite such excess credibility? One familiar suggestion is that numbers carry with them, through their association with the sciences, an aura of authority. I'd like to suggest another mechanism: our use of *cognitive fluency*, a phenomenon well-documented by cognitive psychologists. Cognitive fluency is the "subjective experience of ease or difficulty with which we are able to process information" (Oppenheimer 2008, 237). As it turns out, we often use cognitive fluency as an epistemic heuristic. The easier an idea is for us to comprehend, the more likely we are to accept it as true. This is sometimes a useful shortcut. We are typically better at processing information in domains where we have expertise, so ease-of-comprehension is somewhat correlated with correctness. But the heuristic is far from perfect, as cognitive psychologists

have amply demonstrated. First, we seem more willing to accept an idea simply because it is familiar. Second, we are more likely to accept claims presented in a more legible font. But, obviously, the bare fact of repetition or graphic legibility has no direct bearing on truth. In both cases, using a cognitive fluency heuristic results in a mistaken degree of trust (Reber and Unkelbach, 2010).¹⁸

We should, then, experience a cognitive fluency effect with anything with which are familiar. And we are extremely familiar with numbers. They are the universal abstraction. Information presented in quantified form thus wears an extremely familiar face. So, the fluency heuristic can lead us astray with quantification, just as it does with fonts. This offers an explanation for the unwarranted credibility of quantified values. Fluency may bring somebody to accept a quantified evaluation of value over a more inchoate one — like accepting the *USN&WR*'s clear presentation of a ranking, over one's own internal sense of, say, fit with a law school's culture. And insofar as the quantified presentation is more likely to emerge from external and institutional sources, then the fluency effect gives an unwarranted credibility boost to such sources.

But it's not just that metrics are quantified; it is that they are *standardized*. Once our values are standardized, then we can easily explain our actions and justify our decisions to others. Metrics offer an engineered communicability for values. This engineered communicability grants a further credibility advantage to claims made in the evaluative language of those metrics. After all, our ability to make ourselves understood to others can be a sign that our own understanding is good.¹⁹ And metrics are, by their very nature, easier to understand

¹⁸ I offer a sustained discussion of how cognitive fluency plays into our attraction to seductively clear systems in Nguyen (2021d).

¹⁹ I am relying here on the literature from the philosophy of science's investigation of understanding. According to

across contexts. But there is a gap between communicability and epistemic worth, and that gap can be exploited. As Porter makes clear, institutional metrics trade away informational nuance, richness, and contextual sensitivity for the sake of easy portability across communicative contexts. Metrics, then, in virtue of their basic institutional function, also function to precisely exploit the gap between communicability and epistemic worth. When we standardize metrics, we engineer in broad-ranging comprehensibility by removing contextual nuance. Because the very act of removing contextual nuance increases communicability and cross-contextual comprehensibility, metrics – by their essential nature – invite excess trust.

Quantifications can also be seductive because they offer us the pleasures of *value clarity*. When we internalize them, our value landscape becomes simpler and easier to navigate. We are tempted to take them on, because they offer us hedonistic rewards, in exchange for simplifying our values along certain lines. This line of argument draws on my account of the motivational structure of games, which I've developed at length elsewhere (Nguyen 2019, 2020a). In games we take on artificially constructed goals. In ordinary life, our goals and values are often complex and subtle. It is often hard to explain our values clearly, hard to adjudicate conflicts between values, and hard to figure out if we've actually achieved what we value. But in games, values are easy. They are clearly articulated, with explicit criteria for application. In games, we know exactly what we should be doing, and exactly how well we've done. Games offer us a momentary refuge from the nauseating complexity of real world values. They are an existential balm.

This offers us a second mechanism for the seductiveness of quantification. We can gain a

the standard account, one of the signs of understanding is the ability to communicate that understanding to others (Strevens, 2013). I offer a discussion of how engineered simplicity can hijack our sense of understanding in Nguyen (2021d).

hedonic reward for internalizing simplified values. When we come to value a simplified goal in a non-game activity, we bring the pleasures of value clarity into the real world. Our purposes become clearer, our degree of success becomes more obvious, and our achievements become more readily comprehensible – and it becomes easier to compare and rank our respective achievements. But to get those pleasures, we need to simplify the target. And this helps to explain why it can be so tempting to internalize institutional metrics. Metrics are narrowed and finished. When we internalize such clarified metrics, we can eliminate our struggles with the ambiguity and complexity of our values. Metrics may not have been explicitly made for gamification, but the institutional pressures on the generation of metrics make them function as pleasingly game-like goals.

And value clarity effect becomes even more powerful when that clarity is standardized. After all, the existential burden of our complex values is not merely a personal affair; we have to deal with the buzzing tangle of everybody else's values, too. Navigating this overwhelming plurality – understanding other people's values and explaining our own – can be grueling. There is, so often, a vast gap between our values. Try explaining to another person your profound love of some weird old comedy, or why a sour cabbage casserole makes you feel so comforted on the bleakest of days. Try explaining why a particularly acid passage of Elizabeth Anscombe's fills you with such glee, or why you never quite got along with running, but rock climbing makes you feel so amazing. Sometimes we can make ourselves understood, but often we cannot. So much of our sense of value arises from our particular experiences, the long life we've led, our twisty paths to self-understanding and world-loving — that explaining the whole mess to others is often beyond our capacities.

But institutionalized values offer us an experience of social value clarity. If an institution

offers us a pre-fabricated metric for some value, and we collectively internalize it, then we will make easy sense to each other. Perhaps that pre-fabricated metric is citation rates, or Twitter followers, or GPA, or your university's ranking. In any case, once we internalize that value together, much of the existential friction of social life suddenly disappears. Metrics create a common currency for justification. I no longer need to struggle to explain my way of valuing to others, or to understand their way of caring about the world. Justification becomes easy because metrics can function as a pre-engineered system of aligned value. Metrics offer, not just a personal form of value clarity, but *social value clarity*. When we converge on the same simple, public value system, it becomes so much easier to communicate our values and our justifications. We have gone on the same value standard.

Let's take a step back. Is it some wild accident that institutional metrics turn out to be so seductive? I suspect not, though I can only offer a brief sketch here. Rational agents often need clearly articulated policies to function — including policies about what our goals are, and how we are to evaluate our progress towards those goals. Clearly articulated policies ensure reasonably fast decision making that is consistent over time. As Michael Bratman argues, such policies play an integral part in our being able to maintain coherent agency over time (Bratman, 1987; Holton, 2009; Andreou, 2010). Policies are desirable for large-scale institutions for similar reasons, since institutions also need to ensure relatively quick and consistent decision-making across a large and scattered structure, in order to enable cohesive collective action.²⁰

But the nature of large institutions requires that we heighten the clarity and explicitness

²⁰ This comment relies on the extensive recent literature on group agency, including (List and Pettit, 2011; Gilbert, 2013; Rovane, 2019)

of those policies in order for them to function across consistently across the whole. The policies I set for myself can hinge on my own peculiar sensitivities and ways of understanding the world. A coherent policy for me is: “Exercise every day until I start to get that pleasant warm cheerful feeling.” This works for me because I can consistently recognize that pleasant warm feeling. (Another coherent policy for me is, “Cocktails before 6 PM only when I really, really, really need it.”) But such policies won’t work for large-scale institutions because criteria like “a pleasant warm cheerful feeling” cannot be written into institutional policy, nor could they be reliably applied by different people across the institution. Institutional policies need to be hyper-explicated so that they may be executed by a wide variety of people, hired from a variety of backgrounds. They need to be, to adapt Porter’s language, procedures that are *portable* between many contexts. In order to function in institutions, policies need to be easier to apply – and so they can be appealing to internalize. It is very easy to act clearly and consistently when we adopt such a hyper-explicit policy. However, in adopting them, we are giving up on the kinds of policies that hinge on sensitivity to subtle internal phenomenon.²¹

Context Loss

Let’s step back and summarize the action so far. Values, I have argued, benefit from being tailored to an agent’s particular context. In individual cases, that context can involve all the particular details about the person — their personality, their subtle emotional responses to the world. In group cases, that context can involve details the particular people who make up the group, or the group’s ambient culture. It can involve the kinds of subtle

²¹ For a further discussion of problems with the explicit policies in group agents, see (Nguyen, 2019).

considerations that are only adequately comprehensible to a particular community, who have gone through a particular set of struggles together.²² And in all cases, it involves particular details of the specific context of the agent: the location, the surrounding culture, the environment.

The argument here is an instrumental one, and epistemic in nature. I can easily imagine very different accounts of the harm of value capture. One might wish to argue that autonomous control over one's value was a good in and of itself, and that something was intrinsically wrong with ceding that control via value capture. And this argument might also be a good one. But that is not the argument I am making here. Here, I am arguing that fine control over the expression of one's values is instrumentally good. It promotes well-being and flourishing, and other such ways of indicating a good life, in individuals and communities. It does so because substantively participating in the process of adjusting and fine-tuning values yields more nicely tailored values. Notice this argument works irrespective of whether or not we conceive of the value capture process as consensual or voluntary. It has to do with how much one substantively tailored that value to one's context. One may have voluntarily undergone value capture, but in so doing, one has withdrawn from the process of finely tailoring one's values to fit.

And fine control leads to better well-being and flourishing for epistemic reasons. Individuals and smaller-scale groups have better access to the details of their specific context. There is a useful analogy here to a discussion about the epistemic value of democracy. One reason that democracy is important, one might think, is that self-determination is an intrinsic

²² For an excellent recent overview of standpoint epistemology, see Toole (2022). For a discussion of the tension between standpoint epistemology and the demands of bureaucratic transparency, see Nguyen (2021c).

good. But another reason that democracy is important is that, when appropriately structured, it is the best way to integrate the *epistemic access* of the governed. This is the epistemic defense of democracy. As Helene Landemore (2012) puts it, epistemic democracy functions well when it employs a deliberative process which takes into account the specific details known to the relevant communities. Democratic deliberation, done properly, is sensitive to the special understandings of the deliberating citizens. Importantly, Landemore’s argument isn’t that democratic participation is an intrinsic good, or constitutive of an authoritative government. Landemore’s argument is that democratic participation is instrumentally good, since it yields laws and policies that better fit the circumstances, and take better advantage of the various perspectives, expertises, and understandings of the entire citizenry.²³

The tailoring argument works similarly. When an agent tailors its own values, this yields an instrumental good for epistemic reasons. The agent has more access to their context — their psychology, culture, the local details — and so can tailor their way to a better-fitting value. Value capture involves adopting values from an external source — typically a massive institution. Such an external source has far less fine-grained access to the local details. The large scale at which such institution operates imposes a specific demand: that the information they use can be transferred easily across very different contexts. The general insight from the empirical work on bureaucracies and metrics is this: the larger the scale, the

²³ Interestingly, Landemore doesn’t consider the problems of scale. This is worth a wholly separate discussion, but I can briefly say: Landemore’s argument presumes that the process of democratic deliberation will always preserve the knowledge of the participants, and aggregate them. I think there is significant reason to be skeptical about that. I have offered some reasons to be skeptical in my discussion of the epistemic problems with public transparency metrics (Nguyen, 2021c). I think Landemore underestimates the importance of federalism, and local governance, because of her unwarranted optimism in the possibility for low-loss information aggregation at scale. What the empirical work I’ve discussed on quantification and bureaucracy — especially Scott’s discussion — demonstrates is that information aggregation at scale always leads to massive data-loss. This, Scott suggests, is a reason to strongly prefer local governance in most situations.

less the sensitivity to the details of a particular context. As F.A. Hayek (1945) puts it, central decision makers cannot serve each particular person, but only the average person.²⁴ And, I might add, central decision makers cannot serve local communities, but only the average community.

Value Swamping

You might have started to suspect that there are actually two distinct problems, running side-by-side here: a problem involving externality and a problem involving scale. To disentangle them, let's consider a different phenomenon, right next door to value capture, which will isolate the problems of scale. Consider a case where we actively participate in specifying some shared value – but the efforts of coordination at scale color the formulation of that value. Let's call it value swamping.

Value swamping happens when:

1. An agent's values are rich and subtle (or in the process of developing in that direction).
2. The agent participates in a large-scale social process that yields a specification of shared values.
3. Those specifications of shared values come to dominate the agent's practical reasoning (in the relevant domain).

²⁴ I take Hayek to be a major influence on Scott's analysis. In fact, I think the interest in Scott's analysis is his synthesis of Marxist criticisms of capitalism with Hayekian criticisms of central planning.

Value swamping is just like value capture, except for step 2. In value swamping, the agent doesn't get their values from some wholly external source; instead, they participate in a large-scale social process that yields a specification of a value.

Here is an example, slightly fictionalized from my own life. I took part in a large effort, across the humanities departments of my university, to defend the humanities from constant budgetary incursions from the STEM departments and the business school. We wanted to help the humanities survive and thrive. We ended up in a long discussion about realistic goals, and we decided that we most wanted to push for increasing the number of lines of humanities faculty and increasing faculty diversity. We ended up settling on some targets: we were going to push for a fast increase in the total number of humanities lines by 5%, and embark on a long-term project to increase the representation of people of color in the faculty by 20%.

We needed such clear targets — and such a small number of them — because we needed some specific demands to bring to the upper administration. We also needed highly legible targets — the kind of targets that could be coherently targeted and tracked over the coming years by a revolving set of faculty representatives. Notice what's not on the list, however. I would have loved to push for creative work in new hires, and for diversity in intellectual interests, and not just race. But how could we track such fuzzy, inchoate targets over the years? It might be that every humanities faculty cares about "creativity", but since we lacked a readily accessible and scalable measure of creativity, we can't easily make it a group target.

What happened? Bowker and Star say that any attempt at large-scale collective action creates a demand for cross-contextual informational categories and for data that is readily aggregable. This, in turn, creates a demand for publicly accessible, standardized procedures

of measurement, such as metrics. Notice that the pressure for standardization here doesn't arise from the external sourcing of the metrics, but from the demands of large-scale collective action itself. In other words, for some of the pressures, *it doesn't matter if the metrics are generated by an external source.*

Suppose that our case of value swamping is ideally participatory. Still, as the size of the relevant community scales up, the values that get generated are more and more subject to the demands of cross-contextual communication and consensus. Value swamping admits of more tailoring than value capture. Since a group's values are generated by the group itself, they can still be moderately tailored to the group's experiences. But there are still formidable constraints on the kinds of values the group can use. They can only adopt the sorts of values that can be understood across contexts. The processes of context-stripping and de-nuancing are problems of scale, not of externality. In value swamping, the state-level interests, in information that is aggregable and portable across contexts, aren't the alien impositions are an external force. They are necessitated by the process that we signed onto, and for very good reasons. But they do introduce a pull towards less nuance in the specifications of value, as part of the drive for larger-scale cooperative action.

What we've learned is that there are two problems that can lead to badly tailored values. The first problem is that the values are generated by an external source. The second problem is the values are subject to the pressures of scale. In value capture by institutional metrics, we are exposed to both problems: externality and scale. In value swamping, the values may be our own, but they are still subject to the demands of scale.

Something is still lost in value swamping cases. But we can't quite say that the values aren't our own. We consented, we participated, we actively formulated, and we approve of

the outcome. What's going on here is not exactly outsourcing. But we are sending our values out for processing at a larger scale and getting them back *filtered*. What comes back is what can survive the large-scale deliberative process intact; the more private, intimate, or small-scale communal reasons get filtered out.

The scale problem of values

This suggests a larger picture. What we're starting to expose here is an essential problem with group agency at scale — or at least, a deep tension between smaller-scale agents and the demands of larger-scale agency.

We have lots of reasons to participate in large-scaled collective efforts. Some efforts are far more effectively when scaled up. Some things are best pursued collectively, and sometimes at the largest of scales: reducing carbon emissions, increasing vaccination rates. In many cases, we can pursue those targets most efficiently by agreeing on a precise and shared specification of that target. In those cases, the upsides of having a precise, stable, shared specification of value may outweigh the cost.

But when we scale up our target-setting process, we lose sensitivity, contextual nuance, and granularity. And I take figures like Porter, Scott, and Bowker and Starr to have shown that this is no accident; it is an inevitable cost of scaling up organization for beings like us, using the methodologies we are presently using for informational aggregation. When we need to achieve agreement across a vast scale – across people that don't share all the same context, that don't share the sensitivities – then what we can agree on will need to be subject to that filter of low-context comprehensibility. And insofar as collective efforts require some kind of shared stability, then these collective values will, of necessity, not admit of tailoring

to small scale or individual contexts.

The tensions between the value swamping and contextual tailoring are not the result of some sloppy process of coordination. They are baked into our very nature as limited beings with varying personalities and contexts, who need to coordinate our actions across those contexts. They arise, in particular, from one feature of our finitude: that we each have a special understanding of our own patch and our particular context, and far weaker grasp of distant contexts. So any attempt to render anything comprehensible across scale involves eliminating those details that require special understanding or contextual sensitivity. The tension between small-scale and large-scale valuing is ineliminable; our lives as both individual and social beings will always involve some kind of tension between our small-scale and large-scale commitments. We are beings the kind of beings that are perpetually stuck in a painful compromise – between the intimacy of small-scale understanding, and the de-contextualized comprehensibility demanded of large-scale shared understanding.

None of this shows that we shouldn't scale up our activities sometimes. Some things are best pursued collectively, as shared projects on the largest scales: reducing carbon emissions, increasing vaccination rates. And the demands of large-scaled organization clearly require clear, legible targets. But there is a cost to scaling up. Some goals — like stopping climate change — are worth the cost. And in other cases, we care more about the goods of local tailoring, than the goods of large-scaled cooperation. I can see plenty of good that can come from collectively pursuing a clear target that we all understand in the case of climate change or public vaccination. It is much harder to see the goods that come from collectively pursuing the same specification of, say, values in fitness, or musical values, or values in family organi-

zation. Some things are best managed at a personal or local scale. There is a trade-off, between collective coordination and local specificity – and we may want to make that trade-off quite differently in different domains.

Here is an analogy: in law, we want federalism. That is, we want some of our laws set at the national level, some at the state level, and some at the county or city level. And the explanation is that for some kinds of laws, it's better to coordinate across a vast realm, because the goods of standardization and sameness are worth the cost of low local tailoring. And for other kinds of laws, it's best to let them be set at smaller and smaller scales.

What this suggests is that we should want *value federalism*. Some values are perhaps best pursued at the largest-scale level, some at smaller community levels, and some individually. And the upshot here isn't that we should reject all large-scale values. It's that we should maintain a variety of differently-scaled values. There are many cases in which it might be useful to participate in a larger collective effort, and so to accept, as part of that collective effort, less finely-tailored goals. But, at the same time, we can confine those large-scale, standardized goals to our life inside those collectives, and not let them swamp the rest of our values. The problem occurs when we exhibit an excess preference for the largest-scale values, and let the largest-scale values swamp too many of our smaller-scale values. The problem comes when we let the demand for large-scale legibility intrude into every aspect of our lives, even the most intimate ones.²⁵

²⁵ This paper has had a long, tortured, and meandering history. I owe an enormous debt of gratitude to an enormous number of people, and I am sure that more people have been deeply helpful than I have managed to recall and record. But as a start, I'd like to thank Natalie Ashton, John Basl, Peggy Battin, Sarah Buss, Elizabeth Camp, Elizabeth Cantamalessa, Mercy Corredor, Steven Diggin, Steve Downes, John Dyck, Melinda Fagan, Will Fleischer, Carolina Flores, Rachel Elizabeth Fraser, August Gorman, Joyce Havstad, Max Hayward, Aaron James, Paul Katsafanas, Tim Kenyon, Eliot Michaelson, Elijah Millgram, Paul De Font Reaulx, Ronnie Sandler, Carlos Santana, Alexander Sarch, Karl Schafer, Jonah Schupbach, Nick Stang, Dustin Stokes, Olufemi Taiwo, Greta Turnbull, Elise Woodard, and the reviewers and editors of *JESP*.

Bibliography

- Andreou, C. (2010). Coping with Procrastination. In C. Andreou & M. D. White (Eds.), *The Thief of Time*.
- Arpaly, N. (2004). *Unprincipled Virtue: An Inquiry Into Moral Agency* (1st edition). Oxford University Press.
- Bowker, G. C., & Star, S. L. (2000). *Sorting Things Out: Classification and Its Consequences* (Revised edition). The MIT Press.
- Bratman, M. (1987). *Intention, Plans, and Practical Reason*. Harvard University Press.
- Bratman, M. E. (2014). *Shared Agency: A Planning Theory Of Acting Together* (1st edition). Oxford University Press.
- Brewer, T. (2009). *The Retrieval of Ethics* (1 edition). Oxford University Press.
- Buss, S. (2005). Valuing Autonomy and Respecting Persons: Manipulation, Seduction, and the Basis of Moral Constraints. *Ethics*, 115(2), 195–235. <https://doi.org/10.1086/426304>
- Callard, A. (2018). *Aspiration: The Agency of Becoming*. Oxford University Press.
- Christman, J. (2004). Relational Autonomy, Liberal Individualism, and the Social Constitution of Selves. *Philosophical Studies*, 117(1–2), 143–164. <https://doi.org/10.1023/B:PHIL.0000014532.56866.5c>
-

- Dotson, K. (2014). Conceptualizing Epistemic Oppression. *Social Epistemology*, 28(2), 115–138. <https://doi.org/10.1080/02691728.2013.782585>
- du Gay, P. (2000). *In Praise of Bureaucracy: Weber - Organization - Ethics*. SAGE.
- Elster, J. (1998). *Ulysses and the Sirens: Studies In Rationality And Irrationality* (Revised, Subsequent Edition). Cambridge University Press.
- Espeland, W. N., & Sauder, M. (2016). *Engines of Anxiety: Academic Rankings, Reputation, and Accountability* (1 edition). Russell Sage Foundation.
- Fricke, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Clarendon Press.
- Gabrielle, V. (2018, October 10). How employers have gamified work for maximum profit. *Aeon Magazine*. <https://aeon.co/essays/how-employers-have-gamified-work-for-maximum-profit>
- Gierach, J. 1986. *Trout Bum*. New York: Simon & Schuster.
- Gilbert, M. (2013). *Joint Commitment: How We Make the Social World*. Oup Usa.
- Hayek, F. 1945. "The Use of Knowledge in Society." *The American Economic Review*.
- Hedahl, M., & Huebner, B. (2018). Sharing Values. *Southern Journal of Philosophy*, 56(2), 240–272. <https://doi.org/10.1111/sjp.12275>
- Holton, R. (2009). *Willing, Wanting, Waiting* (1 edition). Oxford University Press.
- Jacobs, J. (1961). *The Life and Death of Great American Cities*. New York: Random House.
- Jaeggi, R. (2014). *Alienation* (F. Neuhouser, F. Neuhouser, & A. E. Smith, Trans.). Columbia University Press.
- Khader, S. J. (2020). The Feminist Case Against Relational Autonomy. *Journal of Moral Philosophy*, 17(5), 499–526. <https://doi.org/10.1163/17455243-20203085>
- Krueger, Joel. (2019). Music as affective scaffolding. In *Music and Consciousness II*, eds. David

- Clarke, Ruth Herbert, and Eric Clarke. Oxford University Press: 48-63.
- Leonelli, Sabine. (2015). "What Counts as Scientific Data? A Relational Framework." *Philosophy of Science* 82, 810–821.
- Leonelli, Sabine. (2016). *Data Centric Biology: A Philosophical Study*. Chicago: The University of Chicago Press.
- List, C., & Pettit, P. (2011). *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford University Press.
- McGonigal, J. (2011). *Reality Is Broken: Why Games Make Us Better and How They Can Change the World* (Reprint edition). Penguin Books.
- Merry, S. E. (2016). *The Seductions of Quantification: Measuring Human Rights, Gender Violence, and Sex Trafficking*. University of Chicago Press.
- Millgram, E. (1997). *Practical Induction* (y First edition edition). Harvard University Press.
- Millgram, E. (2004). "On Being Bored Out of One's Mind." *Proceedings of the Aristotelian Society* 104: 165-186.
- Nguyen, C. T. (2019). "Games and the art of agency." *Philosophical Review*. 128(4), 423-462.
- Nguyen, C. T. (2020a). *Games: Agency as Art*. Oxford: Oxford University Press.
- Nguyen, C. T. (2020b). "The Arts of Action." *Philosopher's Imprint*. 20(14), 1-27.
- Nguyen, C. T. (2021a). "Trust and Sincerity in Art." *Ergo*, 8 (2): 21-53.
- Nguyen, C. T. (2021b). "How Twitter Gamifies Communication." In *Applied Epistemology*, J. Lackey ed. Oxford University Press.
- Nguyen, C.T. (2021c). "Transparency is Surveillance." *Philosophy and Phenomenological Research*. Early Access.
- Nguyen, C.T. (2021d). "The Seductions of Clarity." *Royal Institute of Philosophy Supplements*

89, 227-255.

Nguyen, C. T. (forthcoming). "Value Collapse." Royal Institute of Philosophy Supplements.

Nguyen, C. T. and M. Strohl. (2019). "Cultural appropriation and the intimacy of groups". *Philosophical Studies*. 176 (4): 981-1002.

Oppenheimer, D. M. (2008). The secret life of fluency. *Trends in Cognitive Sciences*, 12(6), 237–241. <https://doi.org/10.1016/j.tics.2008.02.014>

Owens, J., & Cribb, A. (2019). 'My Fitbit Thinks I Can Do Better!' Do Health Promoting Wearable Technologies Support Personal Autonomy? *Philosophy & Technology*, 32(1), 23–38. <https://doi.org/10.1007/s13347-017-0266-2>

Porter, T. (1995). *Trust in Numbers*. Princeton University Press. <https://press.princeton.edu/titles/5653.html>

Reber, R., & Unkelbach, C. (2010). The Epistemic Status of Processing Fluency as Source for Judgments of Truth. *Review of Philosophy and Psychology*, 1(4), 563–581. <https://doi.org/10.1007/s13164-010-0039-7>

Rose, N. S. (1999). *Governing the Soul: The Shaping of the Private Self*. Free Association Books.

Rovane, Carol. (2019). *The Bounds of Agency: An Essay in Revisionary Metaphysics*. Princeton University Press.

Rovane, Carole. (1997). *The Bounds of Agency: An Essay in Revisionary Metaphysics*. Princeton University Press.

Schull, N. D. (2012). *Addiction by Design*. Princeton University Press. <https://press.princeton.edu/books/paperback/9780691160887/addiction-by-design>

Scott, J. C. (1998). *Seeing Like a State: How Certain Schemes to Improve the Human Condition*

Have Failed. Yale University Press.

Sterelney, K. (2010). Minds: extended or scaffolded? *Phenomenology and the Cognitive Sciences* 9 (4): 465-481.

Strevens, M. (2013). No Understanding Without Explanation. *Studies in History and Philosophy of Science Part A*, 44(3), 510–515.

Strevens, M. 2020. *The Knowledge Machine*. Liveright.

Superson, A. (2000). Deformed Desires and Informed Desire Tests. *Hypatia*, 20(4), 109–126.
<https://doi.org/10.1111/j.1527-2001.2005.tb00539.x>

Superson, A. (2010). The Deferential Wife Revisited: Agency and Moral Responsibility. *Hypatia*, 25(2), 253–275. <https://doi.org/10.1111/j.1527-2001.2010.01103.x>

Toole, B. (2022). Demarginalizing Standpoint Epistemology. *Episteme* 19 (1): 47-65.

Velleman, D. (1989). *Practical Reflection*. Princeton University Press.

Velleman, J. D. (2009). *How We Get Along*. Cambridge University Press.

Westlund, A. C. (2003). Selflessness and Responsibility for Self: Is Deference Compatible with
Autonomy? *Philosophical Review*, 112(4), 483–523.
<https://doi.org/10.1215/00318108-112-4-483>