Comments on Seth Lazar's 'Algorithmic Governance and Political Philosophy'

Renée Jorgensen Jan 2023

Seth's lectures deftly highlight the fact that the power wielded by algorithmic intermediaries is significant and deserves careful attention when theorizing about justice. He does this by distinguishing 'extrinsic' from 'intermediary' power. While extrinsic power paradigmatically leverages coercive force, intermediary power can operate more subtly, re-structuring the forms of interaction *available* to the mediatees in the first place. He rightly emphasizes that the tests for whether *extrinsic* power is being justly exercised cannot be assumed to be adequate—indeed, should be expected to be *in*adequate—tests for intermediary power. A model that assumes that the only choices in need of justification are choices to *remove options by threat of violence* is a deeply impoverished model of justice.

1. Distinguishing Power from Governance and Authority

Having identified this power relation, he frames political philosophy's task as answering the questions: what are the right ends, what is the right way, and *who has the right,* to exercise intermediary power? Moving immediately from recognizing power relations to asking about the *authority to wield it* is a very familiar path when theorizing about justice; it is a hallmark of lots of theorizing about *State* power, in particular. But there are other paths to theorizing what makes power relations just, that are (in my view) better suited to the nature of the deeper questions that Lazar aims to excavate. For instance, rather than starting at the top—identifying a *governor* and asking of them *what gives them the right to rule?*—we might start instead with *those over whom the power is exercised*, and ask *what is necessary to avoid their unjust oppression?*

This is the approach (for instance) that underlies Iris Marion Young's 1990 book, *Justice and the Politics of Difference*, and which leads her to characterize the central task for political philosophy as identifying "how institutions and social relationships differentially conspire to restrict the opportunities of some people to develop and exercise their capacities" (Young, 2001:16[1]) or "participat[e] in determining […]the conditions of their actions" (Young 1990:38). Her approach centrally focuses on social structures, "institutional relations and processes of the society" (Young, 2001:2), because it is these, more than the explicitly coercive forces of law, that do the most to shape our lives, protecting us from or exposing us to exploitation, marginalization, and systemic violence; either empowering or rendering us relatively powerless to shape our relations with others.

The criticisms that Lazar raises against theories that focus too narrowly on justifying *coercive* power underscore some virtues of Young's approach. If we started our investigation by asking what opportunities and options the *governed must have* to avoid oppression, then we would naturally meet ***any*** proposed limitation or restructuring of their options by asking what justifies it; we would not

[1] Iris Marion Young, 'Equality of Whom? Social Groups and Judgments of Justice' *Journal of Political Philosophy* (2001)

even be tempted to think (as the criticized theories do) that the question could only apply if *threats of violent coercion* are made.

We would also see that it is a mistake to think that "whenever A has power over B, we should ask whether A has the right to have power over B" (Lazar, 1:17); the implied heuristic simultaneously directs our attention too broadly and too narrowly.

How is it too broad? Because of the entanglements (emotional and otherwise) of our lives and projects, it often happens that A has *power over* B without *governing* them, and thus also without any question arising as to whether A has *authority*, or the *right to govern*, B.[2] How is the heuristic too narrow? Because, much like over-focusing on coercion, taking the issue of whether a power relation is *justified* to be paradigmatically a question of whether *an agent has the right to govern* obscures the many cases where governance occurs—that is, constitutive norms of a social institution are adopted and enforced—but power *over* is *diffuse*. When the unintentionally coordinated actions of many people constrain and enable social groups *differentially*, the fact that the occupants of some social positions are resultantly subject to exploitation, rendered comparatively unfree, or made powerless to participate in decision-making *is* a problem for justice, even when it does not take the form of *an agent* governing without right.[3]

The focus on the *justification of power*—which follows from taking *what gives A the right to rule?* as a starting point—leads Lazar to frame his (otherwise illuminating, important) investigation of the demands of justice on intermediary power as a question about the *all-things-considered justification* of its existence (Lazar 1: 13). This is the wrong question. Consider a similar question with social norms as the object of evaluation. Social norms structure our interactions; they enable and constrain us in various ways, that shape our relationships with each other, our access to resources, our life plans, options, and interpersonal power dynamics. We can ask whether a given social norm is *just*, or good; we can ask whether it leads to *oppression*, or whether we can envision a *better way* to organize ourselves; better norms. But asking whether *the existence of social norms is all-things considered justified* feels like a failure to understand that there isn't an alternative, holding fixed that we continue to interact within social groups. Lazar's replies to the concern that algorithmic intermediary power might be *simply* indefensible have a similar flavor: they aren't so much moral *defenses* as the observation that intermediary power is not optional; we might choose *who* exercises it and *how*, but not whether it exists. This is a good response only if the real question is (as I suggest it is) not *what justifies intermediary power*; but rather *given that there is intermediary power*, how can we make it *just*?

---

[2] If I care *a lot* about Alice's opinion of me, it is descriptively true that Alice has power over me: small words from her will shape my whole deliberative process, engagement with the world, and happiness. She does not have a *positive (claim) right* to exercise this power—I certainly do not have any corresponding duty to submit to it—but her exercise is not *wrongful* either: it neither impinges my rights nor unduly constrains my options.[2] Particularly if we assume (as Seth does) that "the right to govern is […] correlated with a duty not to resist or interfere with that governing power" (Lazar 1:19), it is important to *not* subsume *all power relations* to *governance*. Lazar does not make the mistake of conflating power with governance, but adopts an approach that would make the mistake a hazardously natural one.
[3] These after all constitutes threats both to relational equality and collective self-governance.

2. Has Analytic Political Philosophy Overlooked Intermediary Power?

As my references to Young might hint, I think Lazar is right to urge that we must look beyond justifications of the formal apparatus of coercive state power when theorizing justice for the algorithmic city. But I disagree with his characterization of this as (largely) uncharted territory, and of "political philosophers" as "hav[ing] historically focused on justifying the state's extrinsic power as exercised through law."(Lazar 1:14).

It is true that many prominent philosophers who couch their work in terms of justifying political *authority*—including Raz and Rawls—heavily emphasize the formal state and coercive dimensions of law. But others—famously critical of Rawls' focus as too narrow—pursue the project of theorizing justice as one of avoiding *oppression*. Several prominent feminist political philosophers have contributed to work in this vein, and on *structural injustice* more broadly—including Marilyn Frye[4], Patricia Hill-Collins[5], Susan Okin[6], and Ann Cudd[7]. All emphatically frame issues previously relegated to the 'interpersonal' or 'private' as *political problems* precisely because they structure agents' options (constraining them from the inside out), exercise power *between agents*, and (much like the algorithmic intermediaries that are Lazar's focus) largely constitute our shared "social structure".[8] Their focus is *not* on "justifying the state's extrinsic power as exercised through law", but rather on emphasizing that such a narrow focus leaves out a great deal of power relations relevant to questions of justice; in slogan form, 'the personal is political'. Others theorists (saliently Charles Mills in his 1997 book *The Racial Contract*) centrally analyze the *mediating power of* law—in defining who counts as a person, and hence who can claim the protections of law—as an urgent site for theorizing justice.

In mentioning these scholars, I do not mean to downplay the importance of specifically theorizing intermediary power for algorithmic governance—it is absolutely crucial, for many of the reasons that Seth eloquently identified, to attend it. My aim is rather to point out that our quest to identify the "guiding principles that should shape power relations in the Algorithmic City" (Lazar 1:10) need not start from a blank slate, but rather can and should be informed by the contributions and insights of the related work that has already been done. Precisely because Seth is right that the problem is broader than 'the exercise of coercive power by a central authority', when looking for relevant work, we should cast our net wider than those who pitch their projects as defenses of the authority of the state. The remainder of my comments aim to highlight a few places where more explicit attention to the work of these scholars could have smoothed the path.

---

[4] *The Politics of Reality* (1983)
[5] *Black Feminist Thought* (1990)
[6] *Justice, Gender, and the Family* (1989)
[7] *Analyzing Oppression* (2006)
[8] In a 1987 paper in *Philosophy & Public Affairs*, Susan Okin criticized Rawls' *Theory of Justice* and Walzer's *Spheres of Justice* specifically for failing to recognize that social institutions like *family structure*, interpersonal normative expectations concerning the gendered division of labor, and market structures are all non-state forces that constrain women's options in ways that leave them socially subordinated to men, and so are relevant to justice. She generalized this criticism two years later in her book *Justice, Gender, and the Family.*

3. How Deep are the Asymmetries between Algorithmic and Analogue City?

There are two features of the Algorithmic City that Seth presents as making the issue of just governance particularly acute: First, because algorithmic governance can be dynamically preemptive, it can *constrain* without *coercing*, and moreover can (at least in principle) render it *impossible* to not comply. Second, because (as he quotes Lessig), "Code is never found; it is only ever made, and only ever made by us" (Lazar 1:23:n107) it is not neutral; "every option has to be designed" and specifically enabled.

With respect to preemption, Lazar places algorithmic governance (at least, as imagined in the hypothetically perfected *Pre-Emptopolis*) in stark contrast to law. But the governance possible in Preemptopolis—at least if we understand the example as a city governed by *option design* rather than *literal predetermination*—is not of a *categorically different kind* than is present in analogue cities. To see this, it is crucially important to first recognize a subdivision between what I'll call *formally social options*, on the one hand, and *generic options*, on the other.

A formally social option is *constitutively* formed and *defined by social practices* concerning it. Such an option *in principle* cannot be exercised unless made socially possible.[9] On digital platforms, these might include things like the ability to interact in specific ways with another user's account (attacking their avatar, stealing their digital property, calling them directly in real time, etc.). In the analogue world, these include things like being married, or maintaining credit, as well as scoring a touchdown or winning the World Cup. Formally social options can be made impossible, for all or some, by withholding the prerequisite scaffolding.

By contrast a generic option is an all-purpose option that can be exercised in a variety of ways. These can be 'social' in that they involve others (as buyers, or audiences, or subjects), but are not *constituted* by social recognition. Whether in the analogue or digital realm, it is difficult to enable or disable *just some* instances of a generic option (e.g. prohibiting death threats while permitting speech generally; or prohibiting the sale of specific goods). This distinction cross-cuts modalities of governance, and I suggest that it, rather than the modality of governance, determines the degree of preemptive control available. The illusion of asymmetry arises, I suggest, from focusing on generic options when discussing the limits of law, but on formally social options in illustrating the preemptive power of algorithmic intermediaries.

While recognizing that constitutive laws are a form of intermediary power, Seth suggests (as in discussing the example of marriage in the first lecture) that law is still fundamentally dependent on enforcement; they "can determine whether a given kind of social relation is legally recognized, but cannot (without enforcement) prevent people from participating in the unrecognized counterpart."(Lazar 1:15) The trouble is that many formally social options do not have substitutable "unrecognized counterparts"; and so can be made impossible in analogue spaces just as effectively as in digital ones. And as Young, Mills, and others emphasize, it's hard to see preemptive constraints

---

[9] Lazar calls these 'socially recognized options', but I hesitate to adopt this terminology as it implies the antecedent existence of an option which society then does or doesn't recognize, rather than options that *are constitutively formed* by specific social practices concerning them.

unless they directly affect us; they are often invisible by design (much like the preemptive controls that worry Lazar).

A great deal of legal oppression in this country has been accomplished through omission of this kind. Black Americans were excluded from voting, testifying in court, or standing for election to public office well before the 1860s Jim Crow laws made the exclusion legally explicit. These forms of civic participation were previously simply *not enabled* for Black Americans: they were not defined as eligible to perform these roles.[10] Similarly, well into the 1880s coverture doctrines precluded married women from retaining property in their own name or having personal bank accounts, not by *prohibition* but simply by defining married women as an extension of their husband's legal person. And though credit cards have been commercially common since the 1950s, this option was not "*enabled*" for women—they were not *eligible* to have credit in their own name—until 1974.

We can disobey legal *prohibitions*, and *criminal law* might well have to be public to be effective. But there's a great deal of law that doesn't work this way, and there are many actions that we simply can't do without being recognized as able to do them. Running for public office, building credit, and being married are just some among this set. If not *enabled* by the relevant social or legal structures, formally social options are simply *unavailable.* Whether legal or algorithmic, an agent with the power to withdraw the prerequisite scaffolding "can determine whether they are *literally* possible or impossible."

Generic options are the inverse. We can target undesired instances with bans and content filters, or try to design them away by withholding tools that would make it *easy* to express the banned content (or to manufacture the banned goods)---but all of these measures *are evadable*, by tactics that work just as well in digital as in analogue spaces.

The recent "blank paper protests" in Russia (2022), Hong Kong (2020), and China (2022-3) nicely illustrate the deep similarities between the strategies for evading content filters and moderation in digital space, and those used to circumvent aggressively enforced restrictions on analogue political speech. In the wake of aggressively enforced State suppression of criticism people gathered in public spaces holding *blank* sheets of white paper or poster board, in silence—and were successfully understood as expressing condemnation for both the State's salient recent policies[11], and its willingness to interpret *any* semantic content as justifying repressive punishment. The whitepaper protests in China (or "A4Revolution") were accompanied by an online component: people posted their support for the protest—while evading aggressive algorithmic content filters—by posting walls of text simply repeating the characters for "good", "yes", "correct".[12] Users on Facebook and Instagram signaled support by posting blank white images. Both digital and analogue forms of

---

[10] In 1865, Mississippi and South Carolina passed the first 'Black Codes' (more widely known as 'Jim Crow' laws), formally prohibiting Black Americans from many forms of civic participation. These laws were enforced with the penalty of unpaid labor for white plantation owners. So for the two years until the passage of the 14th Amendment in 1867, Black Americans were *coercively* barred from, rather than defined away from, the exercise of these options. For the rest, they were excluded preemptively, typically by not being granted the status of legal persons.

[11] e.g. the invasion of Ukraine, passage of strict anti-criticism measures, and covid lockdowns, respectively.

[12] https://www.nytimes.com/2022/11/28/world/asia/china-protests-blank-sheets.html

speech were violating anti-criticism regulations from the state, and both successfully evaded the censors' imperfect control.

Of course, an algorithmic content filter can learn, and scrub records of past protest that escaped censorship—but so can the enforcement arm of an authoritarian government. Neither achieves perfect control; in neither case can generic options be "fundamentally limited to what is allowed to us" by our rulers (Lazar, 1:21).

Casting the difference between extrinsic and preemptive power as largely mapping onto the distinction between analogue and algorithmic power misses the underlying truth that both forms of governance exercise huge amounts of *both* forms of power, corresponding to the regulation of formally social options and generic options, respectively. The lesson is not that the preemptive interference Seth points to in the digital realm is unproblematic, but rather that it is also problematic and a source of injustice in its analogue forms. The ability to control which options are enabled—to decide which people, relationships, or actions are *recognized*—is tremendous power *over* and *between*, even (and especially) when it can be exercised without threat of coercive force.

This is not to say that all theorizing work has already been done, or that there are no new challenges posed by algorithmic intermediaries. Some of the features Lazar emphasizes merit special attention concerning the regulation of generic options. The long arm of the state grows substantially longer when it can link traces of individuals across large datasets. And insofar as it is automated, the edicts of algorithmic censors is less dependent on the buy-in of street-level bureaucrats: there are fewer opportunities for discretionary non-enforcement than in analogue equivalents. So some of the ways of resisting increasing authoritarianism are less available within algorithmic governance. Most dangerously, though, we may be more complacent about censorship when enforced by denial of service rather than by imprisonment.[13]

4. Neutrality, Design and Choice Architecture

There is a similar lesson to be learned concerning justificatory neutrality. As Seth frames the issue, it is harder to justify the decisions of algorithmic governance on terms that all can accept, because it must be more opinionated than analogue governance. This involves two claims:

(1) That because "our ability to perform a wide range of the most fundamental actions is not conditioned on the law making those actions possible" (Lazar 1:25), the state can simply establish "a set of starting-gate conditions and get out of the way"(24), taking no view on specific substantive questions.
(2) That because "in the Algorithmic City, the default is that you are *not able to X* unless it is explicitly enabled" (25, emphasis in original), "the only way not to endorse some option is not to enable it, which is expressively equivalent to prohibiting it." (25)

[13] This might embed an overly optimistic picture of our likelihood to object to violent authoritarian censorship. As exercised, persecution for violation of state prohibitions against criticism aren't *advertised* as such; violators simply *disappear*, and often a cover story is offered to create just enough uncertainty that bystanders conclude they lack adequate evidence of flagrant injustice to risk taking a stand against power.

Evaluating these is easier if we consider a concrete example; let's take Twitter. If we define 'option' broadly enough to count *<post text strings that pass a content filter>* as an option, then it is true that every option must be explicitly designed in or out. But equipping users with the ability to exercise such a broad option cannot fairly be read as "endorsing" *every specific purpose* to which they put it (certainly not *endorsing* every meaningful sentence the users assert *in* the text strings.) Once we've fixed a consistent level of specificity for 'option', either it is false that every possible option must be explicitly enabled or disabled, or it is false that the enable/disable choice determines *what one can do with the option* or should be read as *endorsing* the particular exercise. So the second claim looks far too strong.

That said, I agree with Seth that "the appeal to justificatory neutrality seems like a sham if your design decisions precisely enable and constitute differential impacts on different conceptions of the good." (Lazar 1:26). Partly for this reason, we should be wary of the first claim, the idea that analogue governance is more neutral. Lazar writes that for law, "the default is that X is permitted unless it is explicitly prohibited."(Lazar 1:25)—but this is equally true in the algorithmic city. That is to say, it is true *only so long as* we think of permission as merely meaning *not prohibited.* But as my earlier examples illustrate, being denied an *formally social* option through legal omission is no less significant or less in need of justification than explicit prohibition. In letting these cases slip from view, Lazar risks repeating the error for which Young, Okin, Mills, and others criticized Rawls: conflating being "*silent* on specific substantive questions" with "taking no view".(Lazar 1:23)

Moreover, if we consider the decisions made in the analogue city at the same close-range at which Lazar inspects the algorithmic—considering the *administrative* choices, and not only the *guiding principles*—we are immediately confronted with the fact that physical infrastructure is not neutral, either. The spaces of a city—from benches and underpasses to road rules, zoning restrictions, and building codes—must all take some determinate form, and no form will serve everyone's interests equally well. As philosopher Quill Kukla writes (in their 2020 book *City Living)*,

> "There is never any such thing as the unmarked, general way that a space is available to be used or repurposed; different bodies will always fit into and experience a space differently, and be afforded different possibilities for agency within it."[14]

Nothing is truly neutral. The decisions we make concerning a city's physical infrastructure and architecture shape its denizens' prospects, the power relations between them, and do a lot to influence which sorts of social norms are likely to emerge. So it is important to ask, as Kukla does, what it means to say that residents have "a right to the city". They propose that it means the governed—those who live *in* a city—have a claim to "not just access to suitable spaces that meet one's needs and fit one's activities, but also the ability to exercise agency not only within but *upon* these spaces."[15] There are echoes here of Lazar's concern with relational equality, something like autonomy, and self-determination—there is also a promising template for theorizing the sort of participation in self-governance that must be afforded to the governed in algorithmic cities.

---

[14] Kukla, *City Living*, 261.
[15] Kukla, *City Living*, p. 260.

To wrap up: Seth is certainly right that we cannot approach the question of algorithmic governance as rule-from-afar, thinking that if only we select the substantively *right* set of principles, rules, norms, and options, perfected top-down implementation will yield perfected justice. It is a crucial insight that in our digital lives, as in so many other domains, just governance requires (among other things) that we have appropriate *influence* on selecting and augmenting the rules which govern us. Echoing Kukla, we might say we *have a right to the algorithmic city*. But paying more careful attention to the different types of options being regulated shows that rather than being newly urgent or uniquely common in algorithmic governance, the issues Seth raises are *also* a concern in analogue cities.

The problems Lazar calls us to attend to are important problems, and raise new challenges. (I've suggested) the new challenges aren't the degree of preemptive interference available, the modality of governance, or the impossibility of neutral design. What *is* new is the speed and scale of application, and the collective attention to social and relational structures *as designed* and subject to our (collective, imperfect, but effective) control. Theorizing justice for the algorithmic city is an important (indeed mammoth) task; happily, though, we needn't invent it whole-cloth.