Implicit realism impedes progress in psychology: Comment on Fried (2020)

Tal Yarkoni^{1*}

¹Department of Psychology, University of Texas at Austin

*Email: tyarkoni@gmail.com

Fried (in press) argues that progress in the factor and network modeling literatures is currently impeded by inadequate theory development. Here I take issue with this conclusion, focusing on two broad concerns. First, I argue that much of Fried's criticism of previous work (e.g., of general factor models) reflects a particular set of aesthetic preferences and priorities that other researchers are under no obligation to share. Second, I argue that Fried's central argument tacitly assumes a strong realism about psychological constructs that is difficult to defend, and has deleterious practical consequences. When stripped of its realist commitments, Fried's paper provides the reader with little reason to suppose that improved theory development would do much to facilitate progress in psychology. I suggest that applied psychologists may want to consider an alternative possibility—namely, that models constructed at a psychological level of description are simply not very conducive to the production of effective real-world predictions or interventions.

Keywords: psychology, realism, instrumentalism, causality, theory, turtles

Fried (in press; <https://psyarxiv.com/zg84s/>) argues that progress in the factor and network modeling literatures is presently impeded by insufficient or inadequate theory building and theory testing. In this commentary, I argue that there is little in Fried's paper to support this conclusion. I focus on two broad concerns. First, I argue that the central distinction Fried draws between strong and weak (or latent) theories is not well explicated, and implicitly assumes that the reader will share many of Fried's subjective preferences as to what constitutes a satisfactory explanation. While subjectivity is not in itself a bad thing, the implication in this case is that what constitutes a "strong" theory is to an uncomfortable degree in the eye of the beholder: readers whose aesthetic sensibilities diverge from Fried's may find the examples he gives of "strong" theories just as problematic as the examples he gives of putatively weak ones—and conversely, may have perfectly legitimate reasons to favor theories Fried himself dismisses.

Second, I argue that much of Fried's argument tacitly assumes a strong realism about psychological constructs that is very difficult to defend, and has deleterious and largely underappreciated consequences for many psychologists' research practices. I argue that without the implicit dependence on shared

realist assumptions, Fried's paper provides the reader with little reason to think that the kinds of theory-building efforts he advocates for would actually facilitate progress in psychology. I conclude by suggesting that, in the face of perennial pleas for more and better theory in psychology, perhaps the time has come to consider the possibility that what ails much of our field is not too little formal theory, but rather, too great a preoccupation with a certain type of explanatory modeling approach that is simply not conducive to the production of effective real-world predictions or interventions.

One person's description is another's explanation

A central aspect of Fried's argument is the distinction between strong and weak (or latent) theories. Fried defines weak theories as "narrative and imprecise descriptions that are vulnerable to hidden assumptions and unknowns." This contrasts with strong theories, which "explicate a precise set of assumptions and axioms about a phenomenon non-ambiguously". Fried suggests that most theories in psychology are of the former kind: they are specified imprecisely—usually verbally rather than mathematically—and can

be difficult to falsify, as the predictions they issue are ambiguous, and easily reconciled with almost any result via post-hoc auxiliary assumptions.

In principle, there is much to like about Fried’s call for greater precision, formalization, and emphasis on “risky” predictions. Certainly, one would hardly want to argue against these recommendations. What Fried fails to recognize, however, is that many of the weaknesses he identifies in existing models are not weaknesses in any objective sense; rather, they reflect principled trade-offs between values that different researchers may have legitimate disagreements about purely as a matter of intellectual taste. This conflation is intimately tied to Fried’s call for strong theories as a way of arriving at *explanantia* rather than *explananda* (i.e., as causal explanations of phenomena rather than mere descriptions of phenomena that need to be explained). Fried never explicitly tells us what he means by explanation, and he appears to take it for granted that his readers will simply agree with his assessment as to which models are explanatory and which are merely descriptive. I suspect that many will not—and one consequence of this is that the very same criticisms Fried raises against models he dislikes can be used to dismiss models Fried himself favors when deployed by researchers with divergent preferences.

For example, a recurring point in Fried’s paper is his contention that “general factor” models, which posit a single component of intelligence or psychopathology on the basis of factor analyses, are merely descriptive and not explanatory, because “[a] well-fitting factor model cannot be taken as evidence or proof that a psychological construct such as the *p* factor or *g* factor exists as a psychological construct”. Quoting Kovacs & Conway (2019), Fried suggests that such statistical models provide only “...an illusion of explanation: Positing a general factor gives the false impression that there is a psychological explanation, whereas the actual explanation is purely statistical”. But calling general factor models an “illusion” seems rather unfair. For one thing, there is an obvious sense in which *g* and *p* are perfectly viable causal constructs: under the (currently dominant) counterfactual view of causality (Morgan & Winship, 2015; Pearl, 2009;

Rubin, 1974), there is nothing wrong with asserting that individual differences in *g* or *p* are causally responsible for a range of behavioral differences. To say that *g* *causes* differences in lifetime earnings, for example, is simply to say that if we were to directly alter a person’s standing on *g*, holding other things constant, their lifetime earnings would also change. This is a perfectly sensible claim regardless of whether one views *g* as a unitary construct or as a convenient abstraction, and regardless of whether it is easy or difficult (or even impossible) to directly manipulate *g* in practice.

More generally, what constitutes a satisfactory explanation as opposed to a mere description is very much in the eye of the beholder. Fried argues that statistical models of *g* or *p* cannot be considered explanatory models because they have nothing to say about underlying mechanism—that is, they do not explain *why* measures of cognitive ability or psychopathology tend to hang together. But the *why* question implicitly appeals to one’s intellectual tastes. Despite Fried’s intimation to the contrary, many researchers *do* appear to be perfectly satisfied with an explanation of cognitive differences that stops at a high-level, unitary psychometric characterization; these researchers might, if anything, view alternative characterizations—whether they involve more complex network structures or a completely different (e.g., biological) level of analysis—as unnecessarily complicated or excessively focused on implementation details. Conversely, many researchers’ intellectual tastes clearly tend towards lower levels of analysis than Fried’s do, and *those* researchers might very well reject virtually any psychological model—including those Fried points to in his article as positive examples of explanatory models—as a mere description completely divorced from any genuine locus of causality.

Consider, for example, Fried’s section on weak theories in the network literature, where he suggests that “[t]o escape vagaries of language and imprecise predictions, network theory should define what a symptom is, and provide a complete list of all symptoms, causal symptom relations, time-frames, underlying mechanisms, and conditions (e.g. moderators) under which these occur”. That researchers should be explicit

about their causal assumptions is, of course, a welcome recommendation. But the principled approach Fried prescribes here for network theorists seems just as susceptible to the criticisms he raises against general factor models. Why, one might ask, is it any more palatable to treat *insomnia* and *rumination* as causal primitives than to take g or p as such? Is there any doubt that probing the biological or genetic bases of insomnia or rumination would quickly lead to a causal fractionation of these supposedly unitary symptoms, just as with g or p ? Isn't it just as clear that a behaviorally-defined symptom like insomnia is also just an abstraction that emerges from a certain kind of measurement approach, and cannot simply be assumed to "exist as a psychological construct"?

Or consider Robinaugh et al's (2019) elegant computational model of panic disorder, which Fried holds up as an example of how one can set about generating more explicit, precise, and risky predictions. Does the Robinaugh model provide a mechanistic *explanation* of panic disorder, or is it simply a convenient mathematical *description* of observed panic dynamics? The answer depends largely on one's intellectual tastes. To some researchers (presumably including Fried), the Robinaugh model may indeed feel explanatory—it may seem to provide a deeper, more principled way of organizing a morass of disparate behavioral findings. But to many other researchers—e.g., those partial to lower (e.g., biological) levels of analysis—the model is bound to seem hopelessly descriptive and mechanistically underspecified. Would a systems neuroscientist interested in, say, the dissociable contributions of different components of a neural "fear" circuit (e.g., the amygdala, hypothalamus, thalamus, PAG, etc.) consider the Robinaugh model a satisfactory explanation of panic disorder? Doubtful. They would probably say something like, *sure, that's an elegant mathematical model. But it only generates explananda, not explanantia. It doesn't actually tell us why the system produces the dynamics described by the difference equations. For mechanistic explanation, we need to appeal to the underlying biology.*

Of course, our imaginary systems neuroscientist wouldn't want to say something like this within earshot of a molecular neuroscientist, as the latter

could reasonably point out that systems neuroscience is itself sort of a nebulous, underspecified level of description, and one doesn't *really* have a good explanatory handle on panic disorder unless one understands the regularities governing the underlying molecular interactions within the so-called fear circuit. And so on down through chemistry and physics. Once one starts appealing to explanation as way of discriminating good theory from bad theory—or even just theory in general from mere description—it's turtles all the way down¹. Fried argues that, to make progress, researchers will need to "abandon generic common cause and network theories, and embrace the complexity of theoretical and statistical hybrid models", but why stop there? One could just as easily argue that to *really* make progress, researchers will need to leave the factor and network modeling literatures behind entirely, and instead plumb the depths of neurobiology, genetics, or biochemistry.

So far as I can see, there is no principled way to resolve such disagreements, which to a large extent reflect differences in researchers' subjective weightings of values that trade off against each other. For example, animal models generally afford stronger causal inference than observational human studies, but are more difficult to generalize to human populations. Biological models include more mechanistic detail than factor analyses of behavior, but typically have a much narrower scope. A five-factor model of personality is simpler to apprehend and apply than a 300-item model of personality, but, given sufficient training data, will generate poorer predictions (Möttus et al., 2017; Seeboth & Möttus, 2018). Which of these approaches a researcher prefers is only partly a function of objective characteristics; to a large extent, it reflects researchers' differing goals and aesthetic preferences. As Dennett (1991) put it:

When two individuals confront the same data, they may perceive different patterns in them, but since we can have varied inter-

¹I am not suggesting that it's impossible to give a workable definition of explanation, only that any such definition will at some point have to appeal to a subjective judgment about whether or not a given model satisfies the beholder.

ests and perspectives, these differences do not all count as disagreements. Or in any event they should not. If Jones sees pattern α (with $m\%$ noise) and Brown sees pattern β (with $n\%$ noise) there may be no ground for determining that one of them is right and the other wrong. ... Would we prefer an extremely compact pattern description with a high noise ratio or a less compact pattern description with a lower noise ratio? Our decision may depend on how swiftly and reliably we can discern the simple pattern, how dangerous errors are, how much of our resources we can afford to allocate to detection and calculation.”

To be clear, I am not suggesting that the quality of a theory is entirely in the eye of the beholder, and that there is no objective basis to ever prefer one theory over another. On the contrary, many of the concrete problems Fried points to in his paper reflect clear logical or statistical errors that few researchers would try to dispute. Other things being equal, a theory with internally consistent premises is preferable to an inconsistent one; a theory that generates more accurate predictions using fewer parameters is better than one that is both less parsimonious and less accurate; and a theory formalized in a computational model capable of generating precise numerical predictions is better than one expressed purely in vague verbal terms. The trouble is that other things are very rarely equal, and many of the features that Fried takes as impediments to progress in the case of, say, general factor models, are only problematic under one particular set of assumptions and preferences—one that many of Fried’s readers presumably won’t share. Under a different set (one that includes, e.g., a strong preference for compact description, efficiency of measurement, etc.), general factor models will often seem superior to many alternatives—including alternatives that Fried himself favors. And under a still different set (e.g., one with a preference for greater concreteness or biological detail), virtually all of the models Fried discusses in his paper are likely to seem unappetizing.

The general point here is that it is probably not

fruitful to appeal to a distinction between explanatory and descriptive models as a way of justifying the pursuit of “strong” models, unless one also can unpack that distinction in a way that doesn’t depend heavily on subjective judgment—and this, Fried does not do.

Realist intuitions impede progress in psychology

A second theme that runs through Fried’s paper is a strong undercurrent of realism about psychological constructs. By *realism*, I refer to the view that psychological constructs have a mind-independent existence—that is, that at least some constructs have a special ontological status that is not mediated by human thought, language use, or empirical measurement (this view is sometimes called *entity realism*, to distinguish it from other forms of realism, but for brevity’s sake I will use the broader label here). For example, realism about Extraversion implies a belief in some latent quantity each person possesses (to varying degree) that *causes* the various behaviors, thoughts, and feelings that we associate with Extraversion. Any given instrument that purports to measure Extraversion may do a better or worse job at capturing the true construct, but even if the human species had developed *no* good measures of Extraversion, there would still be an objective fact of the matter about whether or not the construct exists, and if it does, to what degree each person has it.

We can contrast the realist position with the view that there’s nothing ontologically special or privileged about psychological constructs like Extraversion, working memory, or self-esteem, and that we invoke and attempt to measure such constructs not because they’re “real” in any special sense, but simply because they *work*—that is, they allow us to generate predictions and intervene on the world in useful ways. What explains the success of scientific theories, if not the actual existence of the posited entities, is a matter of debate². But for present purposes, it

²Perhaps the simplest and most parsimonious antirealist or weak realist explanation for the success of science is that scientific theories *are* capturing real structural regularities in our universe; it’s just that the entities posited by our theories

does not much matter: various flavors of antirealism (Chalmers & Others, 2009; Kukla, 1998), weak realism (Dennett, 1991), pragmatism (James & Burkhardt, 1975), and Fine’s *natural ontological attitude* (Fine, 1984) are all compatible with the core assumption that psychological constructs are not in any sense fundamental features of the universe (i.e., that they are not “natural kinds”), and are always mediated in some way by measurement. I will label this general position *instrumentalism*, as it captures the core idea that we appeal to psychological constructs (both in day-to-day life and in our formal theories) because they are useful to us, and not because they have some privileged ontological status³.

The distinction between realism and instrumentalism is related in an important way to the central psychometric distinction between *reflective* and *formative* constructs (Edwards & Bagozzi, 2000). To a realist, there is a legitimate ontological difference between these two kinds of constructs: when working with reflective models, researchers are postulating the existence of mind-independent psychological entities that have causal force in virtue of their privileged ontological status. When a realist says that Extraversion is a reflective construct, she asserts that each of us possesses some latent, but real, quantity of Extraversion, and that our level of Extraversion causally influences our thoughts, feelings, and actions. By contrast, formative contrasts like socioeconomic status (SES)—where the construct is taken to be caused by the corresponding measurements—imply no such

are only useful approximations of the true data-generating processes—i.e., they’re compressed descriptions of the true causal processes that sacrifice some degree of precision in return for computational tractability.

³Note that instrumentalism as defined here does not require one to give up on the general belief in an objective, mind-independent world beyond our senses; it requires only the narrower belief that our experience of the objective world is mediated through sense organs and faculties whose limitations and biases impose strong constraints on how we carve up that world—e.g., a proclivity to assign simplistic, discrete linguistic labels to statistical regularities that may in reality emerge from very high-dimensional, complex underlying processes. Instrumentalism is thus perfectly compatible with a weak *metaphysical* realism; the relevant contrast is specifically with *entity* realism (Cartwright & McMullin, 1984; Hacking et al., 1983).

commitment.

For the instrumentalist, the situation is different: the distinction between reflective and formative constructs carries no ontological implications. An instrumentalist can have pragmatic reasons for choosing one kind of model over another in a given situation (e.g., because one model is more conducive to making certain kinds of predictions), but that choice carries no implications for the existence or non-existence of any construct. On this view, a statement like “Ellen goes out a lot because she’s highly extraverted” isn’t *wrong*; rather, it’s just a way of summarizing potentially complex statistical regularities in a lossy but useful way, in much the same way that we might say we’re having a particularly rainy summer because of an El Niño event, or that a particular country’s economy is resilient to recession because it’s highly diversified. Are El Niño events and economic diversification “real”? Well, they’re clearly real enough to support useful predictions. But this requires only that a linguistic label like *El Niño* map non-trivially onto some set of statistical regularities in our world’s climate—i.e., that there are, as Dennett (1991) famously put it, *real patterns*. It requires no further ontological commitment to the existence of El Niños as independent causal entities.

Much of Fried’s argument—in keeping with the dominant tradition in the factor and network analysis literatures—is tacitly realist. For example, a good deal of his criticism of “weak” theories focuses on what he argues is a common tendency of researchers to conflate theoretical and statistical models. That *some* conflation of this sort exists and is widespread seems clear, and I stand by the following assertion that Fried approvingly quotes in his paper: “A huge proportion of the verbal claims made in empirical psychology articles turn out to have very little relationship with the statistical quantities they putatively draw their support from” (Yarkoni, 2019). Note, however, that nothing in this quote (or in the paper it is drawn from) has any implications whatsoever for the causal or ontological status of any theoretical construct. The distinction being drawn is strictly between the breadth of the constructs psychologists verbally theorize about and the narrowness of the

statistical operationalizations they employ in their empirical work. The central point Yarkoni (2019) makes is just that it's a mistake to equate a broad verbal construct like "ego depletion" with a p -value from a t -test performed on an experiment contrasting two different conditions of a letter cancellation task (e.g., Baumeister et al., 1998; Charek et al., 2016). This observation carries no causal implications at all, and the mistake in question remains just as much a mistake regardless of whether one thinks ego depletion is a "real" entity or merely a convenient fiction.

Fried, by contrast, views the distinction between theoretical and statistical models through a realist lens: he identifies theoretical constructs with those that exist as real causes of behavior. Much of his criticism of unitary models of intelligence and psychopathology, for example, is based on their putative lack of causal grounding. "Theoretical support," he writes, "...means that there is evidence that g or p exist as something other than statistical parameters." Later, he suggests that "...statistical parameters can be hypothesized to map onto causal processes, but they are not the same kind of thing, much in the same way psychological constructs and latent variables are not the same kind of thing." But the latter distinction implicitly assumes a realist worldview: for there to be any meaningful distinction between psychological constructs and latent variables, one must assume that the former have some privileged ontological status that the latter do not.

Troubles with realism

Unfortunately, this kind of realism suffers from a number of trenchant problems that have been discussed extensively in the philosophical literature (for reviews, see Boyd, 2010; Chakravartty, 2013; Khlentzos, 2011). I will focus on just two here—one metaphysical and one epistemic. The metaphysical problem concerns our understanding of mental causation (or of higher-level causation more generally). Specifically, it's unclear what it could mean to say that psychological constructs are true or real causes of behavior in a world that so far as we understand it seems to be governed by the laws of physics. This problem follows directly from the principle of causal exclusion (Kim,

2007; Malcolm, 1968): if the behavior of all particles is governed by a set of laws accurately described by our best physical theories (currently, quantum mechanics and general relativity), then there would appear to be no room for other kinds of causes. It cannot be the case that my attendance at a party is simultaneously caused both by the evolution of the physical particles that make up my body, in accordance with the laws of physics, and by psychological forces such as my trait level of Extraversion, state feelings of loneliness, etc. Yet this kind of causal overdeterminism seems to be a direct consequence of realism about psychological constructs. It seems lost on most psychometricians that interpreting the arrows in their path diagrams as ontological statements about real causes in the world is not (just) an innocuous move to satisfy folk psychological intuitions, but amounts to a rather profound assault on a conventional physicalist understanding of our universe.

The epistemic problem is that, even if we suppose that some psychological constructs *do* have a privileged ontological status, it's unclear how researchers could ever make that determination in practice—or why it should have any practical implications. How could we definitively know, for example, that the construct of Extraversion seemingly measured by, say, the NEO-PI-R Extraversion scale, is actually real? What if the *real* trait of Extraversion is actually rotated 45° in a 2-dimensional space defined by NEO-PI-R Extraversion and Agreeableness, and personality psychologists have been drawing causal conclusions about a non-existent construct this whole time? (The putative non-reality of measured Extraversion, it should be noted, could pose no barrier to its successful use in practical applications, seeing as any given association between the "true" axes of personality and any other variable could necessarily be equally well captured by linearly combining the axes of infinitely many other parameterizations of the same space.)

Fried, to be fair, explicitly acknowledges this problem of underdetermination. He observes that model-fitting exercises alone can never provide "conclusive information about the data generating (i.e. causal) mechanism". But he fails to explain what (if anything) *could* provide such evidence, suggesting only

that, for a psychological construct to have causal reality, it should map onto some underlying process or mechanism—“...maybe similar to syphilis that was a latent disorder until bacteria in the brains of syphilitic patients were identified in the early 20th century, strongly corroborating the disease as *disease*”. As I discuss later, however, the problem of underdetermination is not restricted to the psychometric level; the mere ability to identify lower-level mechanisms corresponding to a higher-level construct cannot itself constitute a basis for deeming the latter real, since *all* statistically coherent constructs must have a causal basis (after all, they cannot emerge fully-formed from the ether). We are thus still left to wonder: what configuration of empirical results—whether behavioral, anatomical, or molecular—could ever license a psychologist’s conclusion that the *g* factor is not *just* a useful summary of underlying regularities in the true data-generating process, but is actually causally *real*? And if the answer is that there isn’t one—and Fried gives us no reason to think there is one—then the obvious follow-up question is: *why bother with realism at all*? If there’s no principled way to distinguish real constructs from merely useful ones, what do we gain by positing a distinction between latent variables and psychological constructs in the first place?

To be clear, I’m not suggesting that it’s flatly impossible to address these challenges to realism (though I think it’s fair to say that efforts to do so invariably seem to invoke rather convoluted, and at least in my view, wildly implausible, metaphysical assumptions). The point is simply that these are trenchant problems any realist must confront, yet with few exceptions (e.g., Borsboom et al., 2003), psychologists have made no effort to contend with them at all. A psychologist who insists on claiming that the true causes of behavior are things like Extraversion and working memory capacity has the unenviable job of explaining both (a) how psychological constructs could causally influence our thoughts, feelings, and actions without implying a flagrant violation of physical laws, and (b) how one could ever conclusively determine whether a given construct is really real, or is merely a convenient fiction—and why that determination should make any practical difference. An instrumentalist, by contrast,

can readily sidestep both problems simply by asserting that psychological constructs are nothing more than labels we affix to statistical regularities we find useful, and have no reality or causal force independent of the physical laws they ultimately derive from or supervene on.

Some practical implications

The above discussion of realism and its deficiencies might seem like a tangential philosophical digression to many psychologists—*who cares*, some might think, *whether or not the entities posited by psychological theories are really real, as long as they let us intervene on the world in useful ways*? But I think the implicit realism that undergirds a good deal of work in psychology has a number of important, and generally deleterious, practical consequences—in other words, it *does* limit our ability to intervene on the world in useful ways. Here I focus on two particular sets of implications that I argue actively hinder progress in psychology—and particularly the kind of research that presently constitutes a good deal of the factor and network modeling literatures Fried targets.

A bias towards unreasonably simple causal explanations

One way realist intuitions lead psychologists astray is by biasing them towards overly simplistic causal models. If one starts from the assumption that some statistical constructs correspond to real psychological entities while others don’t, then it is very natural to think that having a well-developed psychological theory (i.e., one that articulates which entities exist, and how they interact), should provide a good guide to how one might set about intervening on the world in useful ways. Fried explicitly endorses this coupling between the putative causal status of a variable and its mechanistic basis; for example, in discussing the alleged general factor of psychopathology, he suggests that “...if the *p* factor is indeed a singular causal mechanism that explains comorbidity, the next logical steps may be to find *p*, investigate its neurobiological basis, and try to intervene on it. If, instead, *p* is a sensible and useful summary of the problems people

tend to have ... aiming to identify biological markers for p may make as much sense as trying to find the biological underpinnings for SES.”

Such a conclusion seems intuitive given a realist mindset—after all, what sense would there be in looking for the biological correlates of something that isn’t real?—but there’s actually no sound basis for it. Why should the fact that p is modeled as a unitary causal mechanism in a structural psychological level have any meaningful implications for its realization at the neurobiological level? That p must have *some* neurobiological correlates is not in doubt: since the brain is the proximal source of all human behavior, *all* reliably measured differences in behavior or ability must have neural correlates (for discussion, see Yarkoni, 2015). But this is equally true of constructs we model as formative and conceptualize as merely “sensible and useful summaries”. Suppose we believe p to be nothing more than a useful high-level statistical description of a large set of biological mechanisms that produce pathological behaviors through complex non-linear interactions: why would such a view render the search for biological correlates of p , and for associated interventions, senseless?

The answer is that it wouldn’t. Indeed, it *couldn’t*. If a purely exploratory analysis revealed p to be correlated with, say, the density of oxytocin receptors in the ventral striatum, our chosen psychometric or ontological conceptualization of p could have no possible bearing on whether or not intervening on oxytocin receptors would successfully alter p and reduce “the problems people tend to have”. Reality, after all, is what it is, irrespective of how we choose to think about the measurements we make.

What this kind of thought experiment does reveal, I think, is how pernicious and insidious the downstream consequences of implicit realist intuitions can be. Indeed, in the case of p , Fried’s conclusion (i.e., that investigation of p ’s neurobiological basis only makes sense if p is “real”) appears to directly undermine one of his most important messages—namely, that researchers should strive to embrace complexity. It is difficult to truly embrace complexity, I submit, so long as one is committed to the view that the true

causal processes underlying psychological phenomena must themselves unfold at a psychological level. Researchers who harbor strong realist intuitions are likely to view the search for biological mechanisms or effective interventions as an endeavor secondary in importance to the search for the *true* psychological causes of behavior, under the assumption that identifying the latter will facilitate rapid progress on the former. In reality, there’s little reason to expect any given psychological model, no matter how formal or theoretically pleasing, to provide anything but a very weak—and potentially even utterly unhelpful—set of constraints on the search for lower-level mechanisms.

An analogy may help illustrate the point. Suppose one were able to formulate, on the basis of careful observation of the trajectories of thousands of cars, a mathematical model capable of generating highly accurate predictions of a car’s velocity at time t , conditional on knowledge of its make, brand, and velocity at time $t - 1$. Such a model might be predictively useful in many ways (e.g., it might help a trip computer prospectively anticipate and avoid collisions), but few would expect it to tell us anything terribly useful about how to build a car that runs faster or brakes more efficiently. Why, then, do so many psychologists seem utterly convinced that developing “good” psychological theories of human intelligence or psychopathology is bound to tell us something useful about how to increase cognitive performance or better regulate our emotions?

Do reflective models make sense?

A second set of implications concerns the use of reflective psychometric models. Some commentators have argued that a commitment to realism is a prerequisite for fitting reflective psychometric models. For example, Borsboom and colleagues suggest that “...the standard model in psychological measurement is a reflective model that specifies that the latent variable is more fundamental than the item responses. This implies entity realism about the latent variable ... Maybe more important than this is the fact that psychologists use the model in this spirit” (Borsboom et al., 2003, p. 209). Hood (2013) is even more explicit: “The choice to model covariation with a reflective

model, as opposed to a formative model, requires a commitment to realism. There is no a priori reason to prefer one model to the other unless one harbors realist commitments. Without such commitments, the decision seems arbitrary. Indeed, it would seem irrational to prefer the model that carries with it ontological commitment if one does not embrace some form of realism”⁴ (p. 749).

For what it’s worth, I’m not sure that the above argument is right; I think there could be sensible pragmatic reasons to use reflective models even if one is a hardened instrumentalist and doesn’t believe that the constructs posited by one’s models have any privileged ontological status. For example, suppose we believe that g is merely a useful abstraction, and that what appears unitary when data are factored in a certain way is actually causally explained by variation in hundreds or thousands of distinct biological mechanisms. We might still opt to use a reflective rather than formative model for g if we think that the former model has useful affordances the latter lacks—e.g., that thinking of error as a property of observed indicators rather than of latent constructs will lead us to pay greater attention to the stability of our measurements across different conditions.

That said, even if *some* uses of reflective models can be given a pragmatic justification, it does seem likely that many cannot. There are sizeable research literatures in psychology presently focused on questions that, from an instrumentalist standpoint, seem unlikely to ever lead anywhere productive. Many debates over *the* structure of a psychological phenomenon have this flavor. Does intelligence have a single component, or several? Are there five factors of personality or six? Do “central” symptoms in network models constitute the backbone of depression? It’s unclear that such questions can ever be satisfac-

⁴Note that Hood argues that reflective models may require a commitment only to what he calls *minimal epistemic realism* (i.e., the belief that hypotheses that posit unobservable entities can be justified rationally), and not necessarily to ontological realism. Depending on how one construes this position, it may be entirely compatible with an instrumentalist view, in which case it can be taken as another argument to the effect that reflective modeling *doesn’t* actually require ontological commitments (see the next paragraph of the main text).

torily answered—not because psychology is theory-poor, or lacking in mathematical precision, but for the more fundamental reason that *there is simply no fact of the matter*. The complex, high-dimensional data-generating processes that actually give rise to the variations in thoughts, feelings, and action that psychologists measure in many domains simply do not look anything like the low-dimensional structural models that researchers typically appeal to for causal explanation. This doesn’t mean researchers should avoid such models; they could still have all kinds of important practical uses! What it does imply, however, is that there may ultimately be no objective arbiter of the merit of a theory aside from its demonstrable pragmatic utility. Thus, a researcher would be justified in using a 6-factor model of personality over a 5-factor model just in case the former offers some meaningful practical advantage over the latter, and not because the “true” structure of personality has 6 factors rather than 5—because it clearly has neither (for extended discussion of this particular example, see Wiernik et al., 2020).

Lastly, it’s important to note that even if we take the argument that reflective models require realism at face value, it does not follow that we must embrace realism. As the saying goes, one person’s *modus ponens* is another’s *modus tollens*. Borsboom et al (2003), upon observing that many psychologists rely heavily on reflective models, conclude that realism about psychological constructs is warranted, as it is (at least in their view) required in order for reflective models to make sense. But given the difficulties discussed above, I think it would be at least as reasonable to deny the consequent, and assert that if we believe the realist position to be untenable on logical grounds, then so much the worse for psychologists’ use of reflective models. The fact that the latter conclusion might be unpleasant should not deter us; scientists are not, after all, in the business of deciding what to believe based on what they would like to be true.

Without realism, is there still a case for theory building?

I have argued that psychological realism is a difficult position to defend, and introduces trenchant meta-physical and epistemic problems that psychologists have made almost no effort to contend with. I have also suggested that realism detrimentally affects the way many psychologists approach their research—in particular, by biasing researchers to prioritize low-dimensional structural models they take to map onto “true” causal mechanisms, with little regard for the practical benefits (or lack thereof) of those models. This view is implicit in Fried’s call for theories that ground otherwise “purely statistical” psychological constructs in causal mechanisms, and without it, I think there’s little support in his paper for the claim that a greater emphasis on theory building in the factor and network modeling literatures would facilitate progress. Readers who find themselves sympathetic to Fried’s argument are, I submit, probably sympathetic to it in large part because they tacitly share at least some of Fried’s intellectual tastes and/or realist intuitions—for example, a predilection for network models, or the belief that a defensible distinction can be drawn between latent variables in statistical models and causally “real” psychological constructs.

If one does *not* share such intuitions—and I have argued one probably shouldn’t—then it becomes an open question as to what kinds of methodological or theoretical approaches would maximally facilitate or impede progress⁵. It is possible, of course, that Fried is right, and that the kind of theory-building approach he advocates for would indeed greatly benefit many psychologists if it were more widely adopted. But the paper itself provides little support for this conclusion. Fried does not explain how we could distinguish explanantia from explananda in a way that avoids appealing to researchers’ subjective preferences,

⁵Fried never actually gives a definition of *progress* in his paper, but I assume here that any sensible definition must to a large extent rest on objectively-measurable improvements on real-world outcomes, and cannot simply appeal to one’s subjective sense that things are improving (or, as I discussed at length earlier, the whole thing largely becomes a matter of intellectual taste).

or why the kinds of putatively explanatory models he argues we should pursue should provide a useful guide to successful interventions. He provides no evidence—not even anecdotes—to suggest that the use of formal theories in psychology has historically done much to advance clinical practice, policy making, or other real-world applications in ways that other approaches (e.g., experimental tests of interventions based on researchers’ fuzzy intuitions; atheoretical predictive model development; etc.) have not.

Fried also does not consider the many contradictory data points that might lead us to suspect that perhaps theory-building is actually *not* a very effective way for applied psychologists to make objective progress—for instance, that after over a century of psychology, the first-line treatments for most mental health disorders remain pharmacological interventions developed with little consideration for psychological processes, or that cases where formal psychological theories have been successfully translated into effective public policy remain (at best) rare. He does not discuss biological and genetic evidence that strongly suggests that nearly all major psychological traits and disorders are determined by an enormous array of distinct mechanisms (Boyle et al., 2017; Munafò & Flint, 2014; Yarkoni, 2015)—with the potential implication that most psychological theories, no matter how sophisticated, may simply be too far removed from the underlying data-generating processes to produce practically actionable insights. Nor does he consider ground-breaking advances in machine learning (e.g., the development of deep neural networks capable of human-level or supra-human performance on a wide array of computer vision and language tasks; LeCun et al., 2015) that might lead one to suspect that perhaps the most productive way to make progress on many real-world problems in psychology would be to spend *less* time theorizing about psychological causes, and more time developing extremely powerful learning architectures and training them on massive amounts of data.

To be clear, I’m not asserting that Fried’s claim that lack of theory-building impedes progress in the applied factor and network modeling literatures is *wrong*. I am only suggesting that his paper gives us

little reason to think it's right. It's possible that a primary reason for stagnation in applied psychology is the reliance of many researchers on vague and underspecified psychological theories. It's possible that if researchers were collectively more serious about formal model development, more willing to invest their energies developing *stronger* psychological theories, we would see, just as Fried predicts, a progressive shift away from "generic common cause and network theories" towards more refined theories that generate precise predictions and improve our ability to predict and intervene on real-world outcomes. It's possible.

But I think researchers who got into applied psychology in part because they want their work to have a positive impact on the world would do well to also look around at developments in other fields—for example, in machine learning and artificial intelligence (cf. Yarkoni & Westfall, 2017), where advances in psychology-adjacent domains like computer vision and natural language processing have had enormous real-world impacts in recent years (LeCun et al., 2015; Otter et al., 2020; Sengupta et al., 2020; Voulodimos et al., 2018). They may want to consider the possibilities that a lack of sufficiently strong theory is not *the*, and may not even be *a*, major factor limiting progress in applied psychology; that the challenges the field faces will not be overcome by any strategy—no matter how mathematically or theoretically sophisticated—that prioritizes a search for "real" psychological causes over an ability to deliver practically useful results; that our realist intuitions about how the mind works may have led us astray; and that the kinds of subjectively-satisfying *explanations* we psychologists naturally grasp for may simply not lend themselves very well to the development of effective real-world interventions. That perhaps the root problem is not that applied psychology has too little strong theory, but that it is, if anything, too preoccupied with theory in general—whether good or bad. That now might not be such a bad time to take a little break from the lofty question of *what is causally real*, and to instead try and see if maybe, just maybe, we might get a little farther by pursuing the more pedestrian question of *what seems to work*.

References

- Baumeister, R. F., Bratslavsky, E., Muraven, M., & Tice, D. M. (1998). Ego depletion: Is the active self a limited resource? *Journal of Personality and Social Psychology*, 74(5), 1252–1265.
- Borsboom, D., Mellenbergh, G. J., & van Heerden, J. (2003). The theoretical status of latent variables. *Psychological Review*, 110(2), 203–219.
- Boyd, R. (2010). Scientific realism. *Stanford Encyclopedia of Philosophy*. <https://stanford.library.sydney.edu.au/archives/fall2010/entries/scientific-realism/> Boyle, E. A., Li, Y. I., & Pritchard, J. K. (2017). An Expanded View of Complex Traits: From Polygenic to Omnigenic. *Cell*, 169(7), 1177–1186.
- Cartwright, N., & McMullin, E. (1984). How the Laws of Physics Lie. *American Journal of Physics*, 52(5), 474–476.
- Chakravartty, A. (2013). Scientific realism. <https://philarchive.org/rec/CHASR-7> Chalmers, D., & Others. (2009). Ontological anti-realism. *Metametaphysics: New Essays on the Foundations of Ontology*, 77–129.
- Charek, D. B., Meyer, G. J., & Mihura, J. L. (2016). The Impact of an Ego Depletion Manipulation on Performance-Based and Self-Report Assessment Measures. *Assessment*, 23(5), 637–649.
- Dennett, D. C. (1991). Real Patterns. *The Journal of Philosophy*, 88(1), 27–51. Edwards, J. R., & Bagozzi, R. P. (2000). On the nature and direction of relationships between constructs and measures. *Psychological Methods*, 5(2), 155.
- Fine, A. (1984). The natural ontological attitude. *Philosophy of Science*, 21–44. Hacking, I., Emeritus University Professor Ian Hacking, & Hacking, J. (1983). *Representing and Intervening: Introductory Topics in the Philosophy of Natural Science*. Cambridge University Press.
- Hood, S. B. (2013). Psychological measurement and methodological realism. *Erkenntnis. An International Journal of Analytic Philosophy*, 78(4), 739–761. James, W., & Burkhardt, F. (1975). *Pragmatism*. Harvard University Press.

- Khleutzos, D. (2011). Challenges to metaphysical realism. *Stanford Encyclopedia of Philosophy*. <https://stanford.library.sydney.edu.au/archives/spr2011/entries/realism-sem-challenge/>
- Kim, J. (2007). *Physicalism, or Something Near Enough*. Princeton University Press.
- Kovacs, K., & Conway, A. R. A. (2019). What Is IQ? Life Beyond “General Intelligence.” *Current Directions in Psychological Science*, 28(2), 189–194.
- Kukla, A. (1998). *Studies in Scientific Realism*. Oxford University Press.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
- Malcolm, N. (1968). The Conceivability of Mechanism. *The Philosophical Review*, 77(1), 45–72.
- Morgan, S. L., & Winship, C. (2015). *Counterfactuals and Causal Inference*. Cambridge University Press.
- Möttus, R., Bates, T. C., Condon, D. M., Mroczek, D., & Revelle, W. (2017). Leveraging a more nuanced view of personality: Narrow characteristics predict and explain variance in life outcomes. <https://doi.org/10.31234/osf.io/4q9gv>
- Munafò, M. R., & Flint, J. (2014). The genetic architecture of psychophysiological phenotypes. *Psychophysiology*, 51(12), 1331–1332.
- Otter, D. W., Medina, J. R., & Kalita, J. K. (2020). A Survey of the Usages of Deep Learning for Natural Language Processing. *IEEE Transactions on Neural Networks and Learning Systems*. <https://doi.org/10.1109/TNNLS.2020.2979670>
- Pearl, J. (2009). Causal inference in statistics: An overview. *Statistics Surveys*, 3, 96–146.
- Robinaugh, D., Haslbeck, J., Waldorp, L., Kosakowski, J., Fried, E. I., Millner, A., McNally, R. J., van Nes, E. H., Scheffer, M., Kendler, K. S., & Others. (2019). Advancing the network theory of mental disorders: A computational model of panic disorder. <https://psyarxiv.com/km37w/download>
- Rubin, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5), 688–701.
- Seeboth, A., & Möttus, R. (2018). Successful Explanations Start with Accurate Descriptions: Questionnaire Items as Personality Markers for More Accurate Predictions. *European Journal of Personality*, 32(3), 186–201.
- Sengupta, S., Basak, S., Saikia, P., Paul, S., Tsalavoutis, V., Atiah, F., Ravi, V., & Peters, A. (2020). A review of deep learning with special emphasis on architectures, applications and recent trends. *Knowledge-Based Systems*, 194, 105596.
- Voulodimos, A., Doulamis, N., Doulamis, A., & Protopapadakis, E. (2018). Deep Learning for Computer Vision: A Brief Review. *Computational Intelligence and Neuroscience*, 2018, 7068349.
- Wiernik, B. M., Yarkoni, T., Giordano, C., & Raghavan, M. (2020). Two, five, six, eight (thousand): Time to end the dimension reduction debate! <https://psyarxiv.com/d7jye/download?format=pdf>
- Yarkoni, T. (2015). Neurobiological substrates of personality: A critical overview. In *APA handbook of personality and social psychology, Volume 4: Personality processes and individual differences*. (pp. 61–83). American Psychological Association.
- Yarkoni, T. (2019). The Generalizability Crisis. <https://doi.org/10.31234/osf.io/jqw35>
- Yarkoni, T., & Westfall, J. (2017). Choosing Prediction Over Explanation in Psychology: Lessons From Machine Learning. *Perspectives on Psychological Science: A Journal of the Association for Psychological Science*, 12(6), 1100–1122.