1 The Existential Threats to I-O Psychology Highlighted by Rapid Technological Change

Richard N. Landers

For decades, there has been a quiet murmur of existential discontent within industrial-organizational (I-O) psychology. This has taken many forms, such as calls to mind the science-practice gap (Briner & Rousseau, 2011), expressions of concern over the usefulness of I-O psychology's general approach to science (Highhouse & Zickar, 1997), and calls to increase our influence on and efforts to improve the world at large (Maynard & Ferdman, 2009). Despite decades of commentary encouraging actions to address these concerns, little has changed, and this murmur has in recent years become a bit louder and more insistent, in part because the increasingly rapid pace of technological change, the changing nature of work itself, has made these weaknesses more problematic, more destructive, and more obvious. In short, we are poised to plunge headfirst into our own obsolescence.

In this chapter, my first goal is to explain how we reached this point by describing five key threats to I-O psychology that set us up for this dive. My second goal is to describe some troubling outcomes of these threats so far, to more clearly illustrate why these threats must be addressed. To summarize these outcomes, I-O practice has pulled far ahead of academia in terms of technological expertise, yet in an absolute sense, neither practice nor academia are particularly current or competitive in terms of their understanding of or approach to technology. Third, I provide a list of four recommendations that I believe will turn us toward a better path, one which fully embraces an interdisciplinary future for our field.

1.1 A Perfect Storm for Irrelevance

Some of the threats to I-O psychology I will next describe were created by I-O itself, or more specifically, its culture and common practices, whereas other threats reflect market conditions or the realities of the technological world we now find ourselves in. I will describe these threats in an order of increasing compounding; in other words, each reason is made worse by the reasons that came before it, and in combination, they may be lethal.

1.1.1 Threat #1: Developing Theory for Its Own Sake Is Popular but Not Typically Useful

Numerous I-O researchers over the past decade have noted that I-O psychology literature is becoming more oriented toward an unusual and harmful type of theory development (e.g., Campbell & Wilmot, 2018). To illustrate, consider Table 1.1, which contains a list of titles of articles published in the *Journal of Applied Psychology* from 2018 Issue 1 alongside those published in 1988 Issue 1, thirty years earlier. Even a brief study of this table reveals a noticeable priority shift. Whereas 1988 articles develop measures, investigate effects, and compare methods, 2018 articles are more likely to present theories, test models, and propose mediators. Importantly, my listing of these titles is not to somehow shame or minimize the contributions of either set of researchers or their findings; instead, I use this to illustrate just how abstract and theory-oriented much published I-O psychology research has now become in relation to the I-O psychology of yesteryear. If you have been staying current on the I-O literature, this also should not be at all surprising.

So what might be less obvious to I-Os is that this idea, that the *purpose* of research is to propose theory, puts our field not only in contrast to the historical roots of I-O psychology but also to virtually all research literatures on I-O–related technologies created outside our field. In contrast to I-O theory-building research, technology and the way it is typically researched is highly concrete. In the third column of Table 1.1, I have added a list of recent articles from a respected outlet in the field of human-computer interaction (HCI), an interdisciplinary field that falls at the intersection point between psychology and computer science. In that column, you will find much of the same language of 1988 *JAP*, with lots of measuring, evaluating, and exploring, yet relatively few papers concerning theory as an overarching goal. A cynical traditionalist might interpret this to mean that HCI is 30 years behind I-O, whereas a futurist might interpret it to mean that HCI's increasing popularity must be driven by this applied focus. The truth, as usual, is likely somewhere in the middle. At the very least, this difference reflects a real mismatch between the typical goals of technologists and the typical goals of (publishing) I-O psychologists.

1.1.2 Threat #2: Research on Technology as Yet-More-Stimuli is Artificially Limiting

In the classic language of psychology, technologies are stimuli. They are designed by humans to realize an intended purpose, but once they exist and are in use, they are inherently part of the situations in which people find themselves. People make decisions regarding how to interact with those technologies, or they react as those technologies are forced upon them. Unfortunately, psychology has historically considered and defined its stimuli quite poorly (Gibson, 1960). This is most obvious in social psychology, where even today, stimuli are often developed for use in a single study without extensive pilot testing to ensure that those stimuli are

Table 1.1 *Seven most recent studies across three journals*

JAP 2018, Issue 1	*JAP* 1988, Issue 1	*IJHCS* 2018, Volumes 112–113
Attention to change: A multilevel theory on the process of emergent continuous organizational change.	Development of a new evacuation method for emergencies: Control of collective behavior by emergent small groups.	Head-tracking interfaces on mobile devices: Evaluation using Fitts' law and a new multi-directional corner task for small displays.
A cross-level investigation of informal field-based learning and performance improvements.	Relation of job stressors to affective, health, and performance outcomes: A comparison of multiple data sources.	Evaluating Fitts' law on vibrating touch-screen to improve visual data accessibility for blind users.
Detecting and differentiating the direction of change and intervention effects in randomized trials.	An investigation of sex discrimination in recruiters' evaluations of actual applicants.	A practical approach to measuring user engagement with the refined user engagement scale (UES) and new UES short form.
Cheating under pressure: A self-protection model of workplace cheating behavior.	Effects of preinterview impressions on questioning strategies in same- and opposite-sex employment interviews.	A study of dynamic information display and decision-making in abstract trust games.
The dark side of subjective value in sequential negotiations: The mediating role of pride and anger.	Importance of specialized cognitive function in the selection of military pilots.	Multilingual phrase sampling for text entry evaluations.
On the relative importance of individual-level characteristics and dyadic interaction effects in negotiations: Variance partitioning evidence from a twins study.	Joint relation of experience and ability with job performance: Test of three hypotheses.	Bodily sensation maps: Exploring a new direction for detecting emotions from user self-reported data.
Leadership and member voice in action teams: Test of a dynamic phase model.	Escalation bias in performance appraisals: An unintended consequence of supervisor participation in hiring decisions.	Designing mobile based computational support for low-literate community health workers.

Note. JAP = Journal of Applied Psychology; IJHCS = International Journal of Human-Computer Studies

in fact valid representations of whatever they are intended to represent. This might be attributed to the focus of the field; psychology is, as evidenced by its own name, primarily the study of people's mental states and not the things happening to those people. But such a simple treatment belies the complexity of the world in which people exist. Lewin (1936) already knew this when he stated, "Every psychological

event depends upon the state of the person and at the same time on the environment, although their relative importance is different in different cases" (p. 12). Despite many calls since that time to better integrate both the person and the situation (Ekehammar, 1974), it remains a challenge even today.

When researchers adopt this classic stance, consciously or not, they limit the types of questions that they ask of technology and the approaches they take to studying it. In psychology, such researchers typically default to a stance in which technology takes the form of a well-defined and specific cause, something to either be manipulated by an experimenter or passively recorded in a correlational study, evidenced by research questions like, "Do mobile devices harm measurement?" The reality of technology's relationship with people is more complex, which is recognized explicitly in other fields. For example, in a highly influential article in the field of Management Information Systems, Orlikowski (1992) presented a non-recursive model of workplace technology in which people create and change technology, technology in turn influences organizational policies and norms, and those policies and norms in turn influence how people treat technology; additionally, the technology itself changes how people work, as shown in Figure 1.1. This is a much more flexible and useful approach to studying technology than the simple and uninformative meta-research question "what does technology do to people?" pervasive in psychology and management, the existence of which is in part caused by Threat #1.

Additionally, due to this limited view of technology, specific technologies are often ill defined and misapplied. Grawitch, Winton, Mudigonda, and Buerck (2017) made this argument convincingly and phrased in a way relatable to psychologists: "technology is more than just error" (Grawitch et al., 2017). Importantly, this operationalization of misapplication is not unique to I-O psychology; for example, in media psychology, which is a field that studies the effects of various technologies on human psychology as its primary purpose, researchers still appear to have a significant bias toward investigating psychological concerns instead of technological ones (Reeves, Yeykelis, & Cummings, 2016). In short, because we are trained as psychologists, it is seductive to focus on psychology alone in our research. In the modern world, this approach is often not particularly useful.

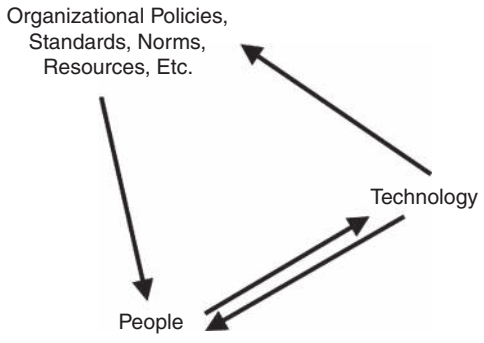

Figure 1.1 *Orlikowski (1992) model of workplace technology*

To remain relevant, we need to be active, integrative, and increasingly interdisciplinary. In contrast to this charge, psychology's mind-set about technology is generally passive, reactive, and siloed. It encourages researchers to sit back and wait until technologies are implemented, often wreaking some degree of havoc upon the world; only when the dust has settled does it become appropriate to begin sifting through what has happened and try to make sense of it. This is, furthermore, reinforced by Threat #1, because one needs to be a passive observer to develop a theory that is only to be tested with confirmatory hypothesis testing, an approach in stark contrast to the natural sciences, where pushing the boundaries of knowledge through invention and discovery are the raison d'être. When is the last time you recall an academic I-O psychologist inventing something new, trustworthy, and immediately useful to practitioners? Although there are a few examples (e.g., De Corte, Sackett, & Lievens, 2011), they are rare, scattered, and tend to fall on the "industrial" side of I-O. It does not need to be this way.

1.1.3 Threat #3: Both Psychology and Technology Are Moving Targets, but Technology Is Worse

The most common epistemology among modern social scientists is likely post-positivism. Many I-O psychologists are not aware of this philosophy of science underlying their research, so I shall take a moment to explore it. Post-positivism, in brief, asserts that there is some "true" state of the world. In statistical terms, these are populations, and within those populations, various relationships, both causal and correlational, are true. So for example, perhaps in the true world, conscientiousness is indeed an emergent state of a person's brain that affects how they behave. We can never know this "true" world; instead, we must make inferences about it via observation, data collection, and statistical tests. Given certain assumptions, we can state with some degree of confidence that our observations in our own world reflect this true world. If I were to stop there, I would be describing the most common philosophical framework behind most modern *natural* sciences, logical positivism. This approach works quite well when measuring the behavior of atoms, or planets, or biological systems, because these relationships are quite stable. The fundamental forces of the universe (i.e., think $E = MC^2$) will not change over time or because we observe them. In psychology as currently studied, this is not a safe assumption. When I conduct a research study to observe the usefulness of Facebook metadata in predicting human behavior, I have no reason to assume between this study and the next that (1) Facebook will be the same, (2) the population using Facebook will be the same, (3) the capabilities of Facebook will be the same, (4) the data being produced by Facebook will be the same, (5) people will behave the same way on Facebook, and so on. Facebook is a living, reactive system, just as the people who use it are themselves complex biological systems. Thus, logistical positivism is not enough for psychology, because (1) researchers need to interpret what they find through these various lenses to make sense of what they find and (2) even if true scores exist, these scores may change over time between one study and the next. Post-positivism is thus a common refinement of

logical positivism that adds these caveats: that we must always reflect upon our own influence, as researchers, on the systems we are researching and also recognize that causal forces from outside the scope of our studies might change the nature of our observations even as we make those observations.

To make this a bit more relatable, realize that post-positivism is the philosophical framework that enables us to conduct meta-analyses of psychological constructs that we explicitly expect to change over time; if we did not believe true scores could move around depending upon when the study was conducted and the assumptions surrounding it at the time, we would expect later meta-analytic estimates to only become more precise, not to fundamentally change. If the true-score relationship between conscientiousness and job performance in 1991 was ρ=.22 (Barrick & Mount, 1991), in a logical positivist framework, we would also expect ρ=.22 in 2091, although measured more precisely. But I suspect most I-O psychologists do not have such an expectation. Jobs will change, people will change, and that number is going to change with them; it is only a matter of how quickly. Thus, even if you have never articulated what post-positivism involves, you probably have an intuitive understanding of it; it is hard-baked into the very foundations of our field.

Why this is critical is that the study of technology on human behavior relies on post-positivism too, although it takes a somewhat different shape. You, as a researcher, do not have the power to personally change the ρ=.22 mentioned above. If the true score is .22 in an organization today, it is very likely to be close to .22 a year or two from now. It may drift over the long term, if the job itself changes, or society changes, or some other "big" thing changes. But it is not something that a researcher, as an individual, can influence. In contrast, modern technologies are constantly being developed, designed, and redesigned by humans according to human needs. Modern technologies are updated continuously with the intent of continuous improvement. Thus, human decisions and behaviors *actively* change true scores between technologies and other variables in ways that are unlikely when examining relationships between psychological constructs alone. If we believe a technology is ineffective in its purpose (i.e., some desirable effect caused by the technology is too weak), we may redesign the technology to increase its effectiveness (i.e., to increase its true score effect). There may be a ceiling to this true effect, given particular design considerations within a particular technology, but there is no clear way to know where either our observed or true scores are in relation to that ceiling.

We have seen the negative effects of assuming technology to be much more stable than it actually is in all areas of I-O psychology where technologies are studied. It is particularly strongly evidenced by the decades-long arguments in our literature regarding assessment center validity (cf., Klimoski & Brickner, 1987; Jackson, Michaelides, Dewberry, & Kim, 2016). The assessment center method, like all selection methods, is a technology, designed by humans to assess other humans' KSAOs. Assessment centers are typically defined by certain common design characteristics, such as the use of multiple raters and exercises (International Taskforce on Assessment Center Guidelines, 2015), but the details vary

dramatically – by purpose, by constructs assessed, by methods employed, by exercises selected, by rater populations sampled, and so on. Thus, as a technology, assessment centers are multidimensional. They incorporate and combine multiple distinct technologies, each with their own quirks, effects, and design considerations. For example, leaderless group discussion is an assessment exercise, and therefore a selection method, and therefore a technology. It can be designed well or designed poorly, and these design considerations are also multidimensional. This logic can similarly be applied to *every technology* contained within *any* assessment center, keeping in mind that some assessment centers may not even overlap with others in terms of the specific technologies employed. This is a startling level of interactive complexity, once the true number of dimensions involved are considered accurately. Furthermore, as the assessment center method has developed, the specific design considerations related to each of these issues have changed; an assessment center designed to the guidelines of 2015 might not have even been referred to as an "assessment center" twenty years earlier. To even *investigate* the "validity of assessment centers" as such in this context is an absolute waste of researcher time and effort. Although the futility of this approach has been recognized to an increasing degree in the last few years (e.g., Kuncel & Sackett, 2014), it took decades to get here. In other technology-oriented literatures within I-O psychology, we face this same road ahead again and again.

As we dig deeper into any technology, whether speaking of the technologies that enable co-located work or the technologies that enable online assessment or the technologies that enable chatbots to teach people leadership skills, the effects of human-contributed variance on true scores will only become more complex. The value of evaluating technologies as if they behave similarly to psychological constructs will remain similarly fruitless. For our field to remain relevant in this new technology-driven landscape, we cannot afford to repeat this same path across every technology-focused research stream within I-O psychology (Landers & Behrend, 2017). This also builds on Threat #2 in that we should not *react* continuously for decades to every innovative technology as it becomes popular, a new stimulus that has appeared suitable for study, pretending that each incarnation of it in our research literature is a random sample from some grand population of technologies. This is unreasonable. And building on Threat #1, neither should we pretend that new technological advancements are simply new versions of technologies we have already studied; our default position should not be to scramble for existing theory as a comfortable and familiar crutch (e.g., Chamorro-Premuzic, Winsborough, Sherman, & Hogan, 2016).

1.1.4 Threat #4: I-O Psychologists Are Not Adequately Trained in Technology

Until recently, it appeared that I-O psychologists, especially those in academia, did not consider technology, as distinct concept needing focused training, to be integral to the field. This is evidenced by Tett, Walser, Brown, Simonet, and Tonidandel's (2013) report on the 2011 SIOP Graduate Program survey, which in part assessed

the degree to which both "substantive" and "methods" topics were covered in I-O psychology programs. Technology did not even make the list of questions, and among what was asked, the most technology-oriented competency area was "human factors." Perhaps unsurprisingly, zero doctoral programs surveyed included this in their curriculum. The next year, Byrne et al. (2014), writing an article inspired by a Society for Industrial and Organizational Psychology (SIOP) panel discussion centered on Tett et al.'s work, described new competency priorities for graduate training in I-O; the word "technology" does not even appear in their work. It is understandable not to *focus* on technology in an I-O psychology graduate program, but this suggests that even just a few years ago, in terms of training new I-Os, technology was not even on the proverbial radar, despite better understanding of technology appearing among the concerns of both I-O students (Harris & Hollman, 2013) and I-O practitioners (Church, 1998; Silzer & Cober, 2010).

Things have certainly changed in the last five years. In 2015, Guzzo, Fink, King, Tonidandel, and Landis (2015) called for I-O psychology to formally respond to the sudden popularity of big data. To inspire I-Os, they provided several examples of I-O work in the big data space already. Yet all their citations to I-O's work in this area appeared in working papers, unpublished manuscripts, and a single published book, all of which were written or published that same year. Importantly, the term "big data" in its current usage has been around since at least 2008, but the concept of analytics at scale had existed for decades before that (Boyd & Crawford, 2012). From this timeline, it is straightforward to conclude that I-O fell a bit behind modern analytics. In response to Guzzo et al.'s article, Aiken and Hanges (2015) called to integrate some degree of modern data science into the core I-O graduate curriculum, including programming skills and modern predictive modeling, primarily suggesting that I-O students should read more books and consider supplementing their own educations by participating in massive online courses on data science until I-O faculty teach themselves enough to in turn teach seminars on the topic. As they noted, "This is not just something that would be nice to see; this is an imperative, and our graduate training needs to reflect this imperative immediately" (p. 544). The threat of technology to I-O became so plain to SIOP that in 2016, the Executive Board established the Future Scanning Task Force to assess threats to the future existence of both SIOP and I-O psychology in general brought by the changing world of work, and to provide recommendations regarding these threats. Understanding technology emerged as a major theme. In 2018, the Executive Board promoted this Task Force to become an Ad Hoc Committee, meaning it will be likely to continue advising the Executive Board for some time. Additionally, two technology-oriented columns intended to teach I-Os about technology now appear in the *Industrial-Organizational Psychologist*: Poeppelman and Sinar's (2016) "The Modern App" and Landers' (2017) "Crash Course in I-O Technology." The push from within for I-O psychologists to understand technology, regardless of application domain, has never been higher.

Despite this increasing pressure, in terms of both initial and continuing education, I-O psychology is struggling to respond. The sudden demand for a new skillset

that most academic I-O psychologists do not have means that there are relatively few people capable of teaching this skillset currently employed to teach graduate students or lead SIOP workshops. This too is changing, although slowly, and Aiken and Hanges' (2015) recommendation to outsource these needs to computer science departments in the interim is unlikely to be successful. Computer scientists have quite dissimilar needs from psychologists in terms of programming expertise, and I-O psychologists are different still. I have chatted with students in I-O graduate programs where this is currently recommended, and, universally, I have heard complaints of perceived relevance and value. I-O psychologists completing programming courses in computer science departments creates the same problem as I-O psychologists completing statistics courses in mathematics departments; it is difficult to understand why what you are learning is useful, and it this kind of contextualization that is presently most critical.

1.1.5 Threat #5: It Is Easier to Bury Our Heads in the Sand

Although this may seem a minor point, it is still worth noting that field momentum is a difficult force to counter. In other words, I-O psychology is a difficult and unwieldy ship to steer. As a field, we are generally decentralized, and SIOP, the European Association of Work and Organizational Psychology (EAWOP), and other national I-O organizations can only do so much. In the case of SIOP, it is a volunteer-run organization, which means that it is in the interests of its leadership to avoid courting controversy. There are no licensure programs or graduate program certification programs to leverage a field-wide shift. Thus, the organization cannot simply tell graduate programs to run themselves differently for the good of the field; instead, committees must be formed, debate the issues, and make recommendations, which the programs can then choose to heed or ignore. This adds significant complexity to decision-making and, more critically, adds a lot of time. I-O psychology, as a field, is about as far from "agile" as is possible, and it is hurting us.

Additionally, finger-pointing is already common. I have heard from numerous I-O academic researchers that this is ultimately the problem of practitioners; academia, after all, can only move so fast. I have also heard from numerous I-O practitioners that the problem is ultimately one of academics; after all, the field has changed, so the training must adapt too. Frankly, neither of these perspectives is productive, as both simply encourage their respective constituencies to "stay the course" on a course that is already off-track. The truth is that I-O psychology, as a field, will live or die together, because these problems are all interconnected (Aguinis, Bradley, & Brodersen, 2014), a so-called "wicked problem" (Behrend & Landers, 2017). The problem with our field's bifurcation is particularly salient in light of Threats #1 – #4. Although practitioners are at the very forefront of exploratory applied research, following and learning about new technologies literally as they change in front of them, it is extremely difficult for any of them to publish in I-O journals given the apparent need to propose novel theory in a confirmatory framework with well-established parameters in every paper.

1.2 Storm Damage So Far

Together, these five threats are interactive; they cause more damage in combination than their individual effects would suggest. This interaction has already manifested itself in at least three ways that promise to become worse if not mitigated soon.

1.2.1 Practitioners Lead the Way in Technology Because Academia Forces Them To

What brought the limitations of academia's approach into greatest relief for me, and really the inspiration for this chapter, come from the results of the first ever SIOP Machine Learning Competition at the SIOP 2018 conference (Putka et al., 2018). In this competition, 17 teams of either academics or practitioners attempted a prediction problem using an authentic turnover dataset provided by a volunteer organization. The dataset was quite large (for I-O research) and complex, with hundreds of variables, systematic missingness, and longitudinal characteristics, among numerous other features. Each team was tasked with creating the predictive model that would hold up the best in a hold-out sample using whatever techniques they had at their disposal. Additionally, teams received feedback on the quality of their models each week for about a month in the form of a leaderboard. Importantly, although academic-practitioner teams were permitted, none formed. At the end of the competition, the top four scoring teams were asked to present on their methods at SIOP. It was revealed that the four winning teams consisted entirely of practitioners.

What is striking about that story, to me, is that academic researchers in both the natural and other social sciences, including the rest of applied psychology, *lead the way.* This is where academics in universities are intended to bring the greatest value, by standing at the forefront of knowledge, unconstrained by organizational politics and the bottom line. Yet, in this competition, the very best minds in machine learning and predictive modeling in I-O psychology were all among practitioners. And perhaps more importantly, very few of the skills used by any of those teams are traditionally taught in I-O psychology programs. Instead, these were all skills picked up in personal professional development, by both the academics and the practitioners, and the practitioners were, as a group, more successful. This suggests that practitioners, or at least academic-practitioner teams, should be leading the charge in our research literature to define best practices and explore the value of all this technology appearing in the employee selection and retention space. So why are there so few such articles? Why are most of the articles we see still building theory of limited practical use?

A troubling truth is that I-O practice, as it is exists right now, is not particularly evidence-based (Briner & Rousseau, 2011). Although this statement prima facie may suggest that practitioners are the problem, the reality is that academia is equally, if not more, to blame. I-O practice does not generally benefit from I-O academia in its current state, because academia is no longer supplying much