

The Flatland Fallacy: Moving Beyond Low Dimensional Thinking

Jolly, Eshin^{*1} and Chang, Luke J.¹

¹ Computational Social Affective Neuroscience Laboratory
Department of Psychological and Brain Science
Dartmouth College
Hanover, NH, 03755

^{*}Corresponding author (eshin.jolly.gr@dartmouth.edu)

Word Count: 5,485

Abstract

Psychology is a complicated science. It has no general axioms or mathematical proofs, is rarely directly observable, and has the privilege of being the only discipline in which the content under investigation (i.e., human psychological phenomena) are the very tools utilized to conduct this investigation. For these reasons, it is easy to be seduced by the idea that our psychological theories, limited by our cognitive capacities, accurately reflect a far more complex landscape. Like the Flatlanders in Edwin Abbot's famous short story (1884), we may be led to believe that the parsimony offered by our *low-dimensional* theories reflects the reality of a much *higher-dimensional* problem. Here we contest that this "Flatland fallacy" leads us to seek out simplified explanations of complex phenomena, limiting our capacity as scientists to build and communicate useful models of human psychology. We suggest that this fallacy can be overcome through (1) the use of quantitative models which force researchers to formalize their theories to overcome this fallacy and (2) improved quantitative training which can build new norms for conducting psychological research.

“Yet I exist in the hope that these memoirs, in some manner, I know not how, may find their way to the minds of humanity in Some Dimension, and may stir up a race of rebels who shall refuse to be confined to limited Dimensionality.” - Edwin A. Abbott, *Flatland: A romance of many dimensions* (1884)

Few works consider the nature of perception and dimensionality as elegantly as Edwin Abbott’s (1884) novella *Flatland: A Romance of Many Dimensions*. The narrator of the story, A. Square, lives in a world full of “Flatlanders” incapable of perceiving or even conceiving of a reality that exists beyond two-dimensions. However, after a visit from a “Stranger” (a sphere) A. Square comes to appreciate how complex and high dimensional the world really is. Ultimately, he is imprisoned for his heretical beliefs after trying to teach his colleagues about his revelations. A key insight from Abbott’s work is that creatures with limited perceptual capacities (i.e. seeing in only two dimensions) come to *reason* in a limited way, ignoring the complexity of the world and truly believing their perceptions to be veridical (Figure 1A). Much like Flatlanders, humans exhibit strong biases in their reasoning about a complex and high-dimensional world due to finite limitations on their cognitive capacities. For this reason, we posit that psychological researchers, constrained by these limitations, may come to view the complexities of psychological life in similarly limited ways.

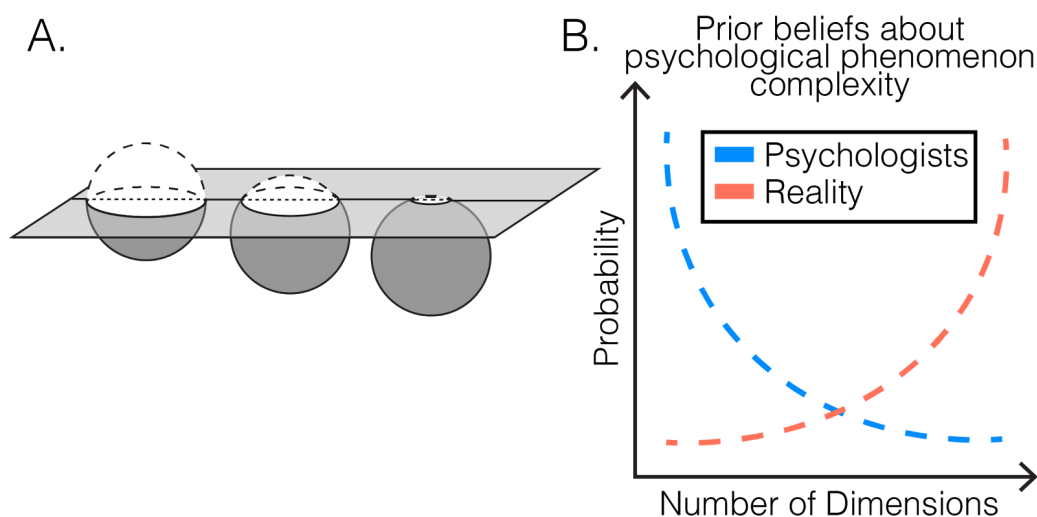


Figure 1 | A) In Abbott’s novella, A. Square cannot perceive his world as anything other than two dimensional. From his limited perspective, a three-dimensional entity (sphere) appears to be changing shapes before him (growing and shrinking circle). In reality, this entity is simply moving through a lower-dimensional plane, but A. Square’s limited perspective leads to a false conclusion about the nature of reality. B) For similar reasons, psychological scientists may falsely conclude that the number of dimensions that accurately characterize psychological phenomena is sufficiently small, even if in reality the complexity of psychological life is necessarily high dimensional. This toy diagram illustrates the probability distribution function reflecting how psychological researchers may have strong prior beliefs (indicated by high probability values) about the complexity of psychological phenomena.

Human psychology is rife with complexity, the product of an immensely high-dimensional space characterized by interactions between trillions of neural connections, billions of unique individuals, and dynamic changing contexts spanning thousands of years of history. Despite this complexity, the majority of theoretical developments in psychological research have consistently converged on producing a number of low (typically two) dimensional/factor theories/process models of human mental life. Since at least the early work of Plato (Evans & Frankish, 2009) this manner of characterizing the mind has come to dominate our understanding of *emotion* (Damasio, 1994; Davidson, 1993; Ochsner, Bunge, Gross, & Gabrieli, 2002; Posner, Russell, & Peterson, 2005; Russell, 1980; Schachter & Singer, 1962; Zajonc, 1980), *social cognition* (Chaiken & Trope, 1999; Cuddy, Fiske, & Glick, 2008; H. M. Gray, Gray, & Wegner, 2007; K. Gray & Wegner, 2009; Greenwald & Banaji, 1995; Haslam, 2006; Mitchell, 2005; Saxe, 2005; Todorov, Said, Engell, & Oosterhof, 2008; Waytz & Mitchell, 2011), *moral judgment* (Greene, Sommerville, Nystrom, Darley, & Cohen, 2001; Rand, Greene, & Nowak, 2012), *learning* (Daw, Niv, & Dayan, 2005; Frank, Cohen, & Sanfey, 2009; R. A. Poldrack et al., 2001; Russell A. Poldrack & Packard, 2003), *cognitive control* (Heatherton & Wagner, 2011; Hofmann, Friese, & Strack, 2009; McClure, Laibson, Loewenstein, & Cohen, 2004; Schneider & Shiffrin, 1977; Shiffrin & Schneider, 1984), *decision-making* (Chang & Sanfey, 2008; Dijksterhuis, Bos, Nordgren, & van Baaren, 2006; Kahneman, 2003; Sanfey & Chang, 2008; Sloman, 1996; Wilson & Schooler, 1991) and *reasoning* (Epstein, 1994; Evans, 2003; Stanovich & West, 2000). How could something as complex as the human mind be consistently described in two dimensions, irrespective of the mental faculty under consideration (Figure 1B)? Although these theories have provided a bedrock for empirical investigation, we argue that rather than reflecting a rich characterization of the complexity of human psychology, they instead reflect a simplistic view of our scientific understanding (Flatland fallacy)—a product of the limits of our cognition.

In this paper, we outline several reasons why we believe psychologists consistently converge on two-factor solutions to characterize our understanding of human psychology. We argue that these conclusions arise from our limited cognitive capacities, social norms ubiquitous in field of psychology, and our reliance on language and low-dimensional visualizations to communicate research findings. We suggest that moving beyond low-dimensional thinking requires formalizing psychological theories as *quantitative* computational models capable of making precise predictions about cognition and/or behavior, and advocate for improving training in technical skills and quantitative reasoning in psychology.

Why does the Flatland fallacy happen?

Understanding why the Flatland fallacy occurs requires examining both biases and limitations in human cognition as well as cultural norms in research and training in psychology. Specifically, we propose four main reasons why this fallacy occurs: (a) biases in the feeling of understanding; (b) limitations of human cognition; (c) over-reliance on traditional experimental design and analytic approaches; and (d) limitations in our ability to communicate complex concepts. Because psychology researchers have the unique privilege of being members of both the matter of study and those conducting the study, it is critical that the products of our science not be constrained by the limits of our own psychology (Meehl, 1954).

Feelings of Understanding

Much like Abbott's A. Square feels that a two-dimensional existence is a complete account of his universe, humans are prone to a "folk understanding bias"—the sensation that simplistic explanations lead us to believe we truly understand more complex phenomena. Prior work in cognitive science and philosophy has illustrated how individuals can fall victim to cognitive biases that lead them to believe their actual understanding of phenomena exceed their true understanding of phenomena. For example, individuals often report a high feeling-of-knowing despite their inability to accurately recall previously learned information (Koriat, 1993). When individuals are tested on their ability to explain how a system works (e.g. a quartz watch) they tend to report an overestimate of their knowledge until they are asked to provide a specific explication (Rozenblit & Keil, 2002). In other words, individuals create mental placeholders of elaborate, in-depth explanations (e.g. essences and hidden mechanisms) that give rise to a feeling of certainty and understanding, even when limited understanding exists (Medin, 1989; Strevens, 2000). Because these approximations can provide basic explanations for how a system works, they are initially insightful, leaving people with the sensation that they know more than they really do (Rozenblit & Keil, 2002).

A critical observation is that this bias is exacerbated when individuals are asked to explain systems that are highly opaque, i.e. have poor observability of their inner workings (Rozenblit & Keil, 2002). Unfortunately, the complexity of psychological science lies almost entirely in its lack of transparency; mental processes are not directly observable, only inferrable through observations of behavior and their correlations with biological functioning.

This bias likely originates from our strong motivation to understand and find meaning in our experiences and the world as a whole (Cohen, Stotland, & Wolfe, 1955). For this reason, it makes sense to favor simplistic and easily understandable theoretical conclusions over complex and complete accounts of phenomena, even if they are only weakly supported by experimental data. Simple explanations provide some uncertainty resolution even if they paint an incomplete picture (Pinker, 1999; Webster & Kruglanski, 1994). Consequently, researchers may be collectively at risk for pursuing a psychological science that they can "understand," irrespective of whether that science offers robust predictive accuracy.

Limitations of Cognitive Capacity

Limitations in the cognitive capacities and motivations of individuals offer another explanation for the Flatland fallacy. It is well established, for example, that humans are not supercomputers who always calculate mathematically optimal solutions for the problems they face (Gigerenzer & Goldstein, 1996). Rather, our brains are the product of specific evolutionary constraints such as physical *size*—they need to be small enough to permit live births and allow us to locomote; *speed*—they need to support processing that can occur on finite time scales; and *energy*—their energy demands cannot exceed our metabolic abilities (Montague, 2007). For this reason, the notion of bounded rationality has been instrumental in characterizing how the mind processes information quickly and reasonably accurately (Simon, 1957). There are many well known examples, such as our limited capacity to simultaneously manipulate large chunks of information (Miller, 1956), process multiple attentional tasks (Pashler, 1994), and

uniquely represent person-identity information without relying on feature similarity such as stereotypes (Mervis & Rosch, 1981; Smith & Zarate, 1992).

We believe that the key psychological limitation underlying the Flatland fallacy is our inability to reason in more than a few dimensions, particularly in contexts that require integrating multiple sources of information together. Individuals tend to default to simpler, general, heuristic-like strategies that serve to make such reasoning more cognitively tractable. These strategies often constitute lower-dimensional approximations (e.g. two or three) of far more complex information landscapes, which raises the possibility that even the process of conducting scientific research can be similarly marred by the limits of our cognition. We outline three examples of how lower dimensional approximations impact how we make judgments and decisions.

Judgment. Many real-world settings involve situations in which individuals are faced with the task of making judgments by combining a large number of potentially relevant factors (e.g. clinical evaluations). A large body of work has consistently demonstrated that humans make judgments using just a handful of dimensions (Brunswick, 1952) rather than considering all the available information on hand. In particular, individuals overweight or underweight the relative importance of specific factors or simply ignore seemingly irrelevant information in favor of simplified evaluation criteria, such as when estimating school admissions, personality metrics, and even criminal evaluations (Dudycha & Naylor, 1966; Karelaia & Hogarth, 2008; Meehl, 1954). In other words, individuals rely on heuristics to simplify the space of information under consideration, especially when this space is very large or shares a nonlinear relationship with an outcome (Deane, Hammond, & Summers, 1972; Karelaia & Hogarth, 2008). Given the robustness of this evidence, there has been a strong call to incorporate statistical models that can integrate more dimensions in place of solely relying on clinical judgment to overcome these cognitive limitations (Dawes, 1971; Dawes, Faust, & Meehl, 1989; Meehl, 1954).

Decision-Making. Dual-process theories have a rich history of characterizing human decision-making (Sanfey & Chang, 2008). These accounts have been used to explain how we can switch between fast, intuitive, and reflexive modes of thinking to slow deliberative calculations (Kahneman, 2003, 2011) and also how emotions and cognitive deliberation might be integrated when making decisions (Chang & Sanfey, 2008; Chang, Smith, Dufwenberg, & Sanfey, 2011; Greene, Nystrom, Engell, Darley, & Cohen, 2004). In addition, there appears to be an upper limit on the number of attributes that can be simultaneously considered when making a decision between different choices (Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Payne, 1976). Faced with a large number of factors to consider, individuals appear to act in an adaptive manner falling back on heuristic shortcuts rather than considering all of the available information on hand (Payne, Bettman, & Johnson, 1993; Simon, 1987). Specifically, as the number of alternatives or attributes under consideration increases (e.g. the number of apartments an individual is choosing amongst), the number of dimensions that individuals ultimately utilize to make a decision between the alternatives decreases (Timmermans, 1993). In other words, humans tend to make high-dimensional problems more cognitively tractable by

considering lower dimensional perspectives — specifically, through the use of heuristic strategies that entail ignoring potentially relevant factors (Gigerenzer & Brighton, 2009).

Conditional Reasoning. Even in more socially interactive contexts that require individuals to consider the motivations of *others*, individuals exhibit consistent limits on their cognitive abilities. A large body of work in game theory and behavioral economics has demonstrated that individuals are limited in their depth of strategic reasoning (Camerer, Ho, & Chong, 2015). In this work, individuals compete or coordinate with each other within an economic game. Players' strategies in these games allow for optimizing their own payoffs while also considering the strategy utilized by other players. Interestingly, individuals are rarely able to reason more than two steps ahead of other individuals, (i.e. more than two levels of such conditional reasoning: "I think that you think that I think") (Camerer, Ho, & Chong, 2004; Griessinger & Coricelli, 2015; Stahl & Wilson, 1995). Even in non-strategic contexts, individuals have been shown to exhibit limits on the amount of recursive reasoning they are capable of, such as during theory-of-mind tasks that require inferring the motives of fictionalized agents (e.g (Happé, 1994)) or more generally, parsing language comprised of numerous embedded clauses (Karlsson, 2010). Across a large number of spoken languages, for example, this type of syntactic recursion (e.g. "a car the man the woman the boy saw drove fast") rarely exceeds two levels of depth and even in written text rarely exceeds three levels of depth (Karlsson, 2007). Because understanding conditional complexity quickly becomes incredibly difficult, individuals fall back on using simple heuristic strategies and generalized decisions rules in lieu of making more optimally rational choices (Camerer, Johnson, Rymon, & Sen, 1993; Costa-Gomes, Crawford, & Broseta, 2001).

Taken together, these findings suggest that in the face of complex information processing, individuals intuitively converge on strategies that reduce the number of factors (dimensions) under consideration to make cognitive problems more tractable. Because the dimensionality of factors necessary for understanding human psychology is incredibly high, psychological researchers may be focusing on far fewer dimensions than what actually comprise the mind. That is, scientists may intuitively converge on establishing low-dimensional theories (e.g. dual process models) because they allow for a reduction of the diverse set of factors relevant to building a comprehensive account of psychological processing.

Cultural Norms

Together with individual biases and cognitive limitations, we believe that methodological traditions within the field of psychology have built a pedagogy that supports the Flatland fallacy. Heavily inspired by Ronald Fisher's iconic *Statistic Methods for Research Workers* (Fisher, 1925) and deeply embedded in undergraduate training in psychology, there is a strong persistent cultural tradition of academic psychology's reliance on a specific type of experimental design and statistical analysis to test hypotheses: two dimensional factorial designs (i.e. two-way analysis of variance; 2x2 ANOVA) evaluated via null hypothesis significance testing (NHST).

Although psychology does not have a generally agreed upon core curriculum, at minimum almost all psychology undergraduate programs require that students take an introductory psychology course and one or more courses in statistics or research methods

(Stoloff et al., 2009). Introductory statistics courses typically introduce basic concepts of inferential statistics, culminating in an introduction to two-way ANOVA. Research methods courses primarily emphasize making causal inferences using 2 x 2 factorial designs. In addition, approximately 40% of psychology programs require taking 5 courses in specific topics of psychology (e.g., abnormal, developmental, cognitive, social, biological) and a capstone course that results in a culminating experience for the psychology major (Stoloff et al., 2009). We believe this pedagogy not only trains generations of psychological scientists to pursue empirical investigations that favor simple factorial designs, but also teaches them to *think* about psychological science in a low-dimensional way. Consequently, the standard approach for inferential reasoning in psychology produces theoretical ideas that are limited to measuring mean differences between variables manipulated along a handful of psychological dimensions, thereby perpetuating the Flatland fallacy.

Although this approach provides an accessible starting point for the evaluation of psychological phenomena, significance testing of experimental manipulations alone do not constitute a formal model of human psychology and cannot be used as evidence for such (Bolles, 1962). The pervasive danger of relying on this approach when testing psychological theory is that it reinforces the generation of weaker, less specific and more nebulous theories, particularly in a growing era of “big data” (Meehl, 1990; Van Horn & Toga, 2014). This methodological paradox was elegantly articulated 50 years ago by Paul Meehl (Meehl, 1967). Whereas increasing experimental power constitutes a *more difficult* test that a quantitative theory must pass to remain viable (e.g. a mathematical model in physics), the exact opposite is true of psychological research, which is primarily concerned with utilizing NHST to detect arbitrary non-zero differences between experimental conditions. Because measurement error decreases with increasing power and precision, smaller non-zero differences will be *necessarily* more detectable in the limit of NHST. This leads psychological researchers to conclude that a statistically significant result of a trivially small difference provides support for a given theory. By instead developing theories as the “prediction of a *form* of function (with parameters to be fitted)” or “prediction of a quantitative magnitude (point-value)” the band of tolerance around theoretical validity *decreases* as experimental fidelity increases (Meehl, 1967). In other words, using NHST to evaluate theories formalized as models that make specific quantitative predictions ensures that theories must be *specific* in order to remain robust.

A classic example of the benefit of such an approach in psychology can be found in prospect theory (Kahneman & Tversky, 1979). Prospect theory outlines the form of a function that describes how the quantitative gain or loss that an individual incurs is mathematically transformed into the subjective *value* they feel (Figure 2A). Critically, this theory defines a weighting function (Figure 2B) for outcomes that can not only be used to predict individuals’ decisions, but also captures the asymmetry that occurs between changes in income framed as gains or losses. This model differs drastically from conventional approaches discussed previously in that it can be used to make point predictions about individuals’ behavior and attitudes in a wide variety of decision contexts (Tversky & Kahneman, 1974).

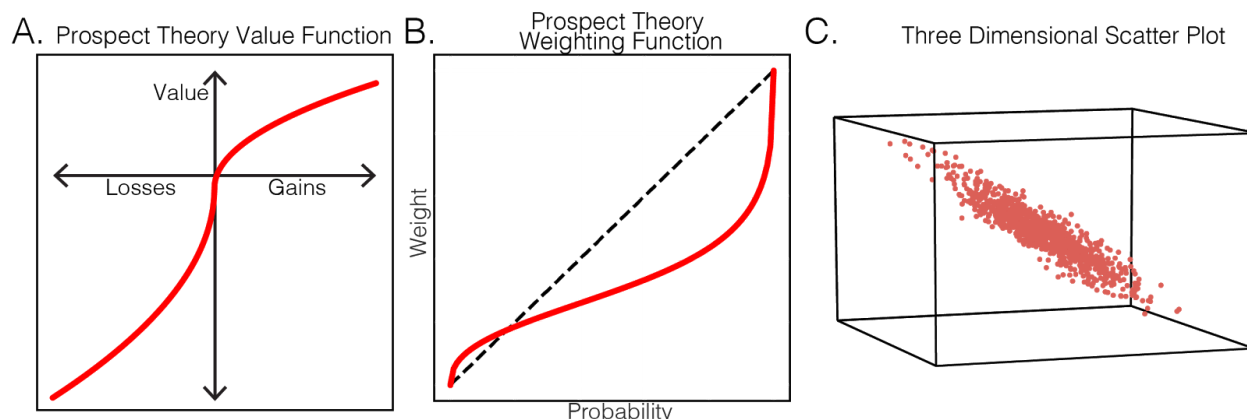


Figure 2 | A) Prospect theory describes a mathematical *function* that maps between financial gains and losses and the subjective value that individuals experience. This function can be used to predict how an individual will make decisions and explains their tendency to exhibit risk-seeking or risk-averse behavior. B) A central part of this theory is a functional account of how individuals treat probability values when making decisions, illustrating how small probability events receive more consideration than large probability events when making choices. C) Researchers are always limited to a small number of dimensions to graphically depict their findings. Even with advanced visualization tools utilizing dimensionality reduction techniques like principal components analysis (e.g. (Heusser, Ziman, Owen, & Manning, 2017)), care must be taken to interpret such visualizations as low-dimensional *approximations*, limited in their ability to fully characterizing high-dimensional landscapes.

Communicating Complexity

Lastly, we consider how a pervasive communication problem endemic to the field of psychology as a whole gives rise to the flatland fallacy. Central to this issue is that collectively, psychological researchers lack a *lingua franca* that enables us to communicate about the complexities of our field. While numerous succinct psychological constructs exist to describe complicated ideas (e.g. “stereotype threat” (Steele & Aronson, 1995), “affective forecasting” (Gilbert & Wilson, 2000)), these ideas are not supported by a common and extensible framework for communication and development. In contrast to social science fields that rely on verbal descriptions of phenomena (e.g., psychology) (Watts, 2017), scientific disciplines that have methods of communicating about high-dimensional problems using mathematics (e.g., physics, computer science, economics), concrete physical models (e.g., biology, astronomy), or their own notation (e.g., chemistry) may be less susceptible to the Flatland fallacy. We argue that in the absence of a foundational dialect, psychology has suffered from an inability to communicate complexity, instead generating non-predictive, ad-hoc hypotheses *after* the observation of data (Kerr, 1998; Meehl, 1990). As a result, psychological discourse is more like a pidgin language than a lingua franca: simplified mixtures and non-specific generalizations that make it difficult to build a cumulative science.

This communication problem is particularly exacerbated by the limited number of dimensions available to psychologists to visualize their findings. In theory, visualizations provide a useful tool for detecting patterns and interpreting research findings. However, even with compelling visualization techniques such as hypertools (Heusser et al., 2017) and t-SNE (van

der Maaten & Hinton, 2008) (Figure 2C), graphical representations of research findings are typically limited to about three dimensions¹. This is particularly problematic in the absence of a formal framework to build a cumulative science (e.g. mathematics), as psychologists are only able to interpret experiments that independently manipulate a few dimensions (e.g. a three-way ANOVA). This creates a tension between theoretical interpretability and theoretical extensibility. Because psychological science lacks a formal discourse, psychological scientists are motivated to design experiments around low-dimensional theory testing and simple visualizations, but because experiments designed to test low-dimensional theories are not by themselves extensible, psychological science exists as a patchwork of disparate, loosely connected ideas. Even recent advances in “theory mapping” (K. Gray, 2017), which provide organizational instructions for how to connect psychological theories to each other, fail to provide comprehensive *models* of psychology that can be used to make useful quantitative predictions (Yarkoni & Westfall, 2017) about behavior or cognition.

What are some solutions?

We believe that overcoming the Flatland fallacy requires advances in the methodological approach that psychological scientists take towards their own work and also pedagogical changes that train future generations of researchers to build upon and extend extant work. First, we highlight the role that *computational models* can play in enabling researchers to overcome the biases and limitations in their own cognition, as well as enabling multiple researchers to work together to build a more cumulative science. Second, we outline some suggestions for the improvement of *psychological training*, stressing the importance of teaching students the fundamentals of mathematics and computer programming. We believe that both of these approaches are required to mature the field of psychological science.

Computational Models

Formalizing psychological theories using computational models provides a way to overcome the Flatland fallacy through the consideration of *high dimensional* explanations of psychological phenomena. Indeed recent methodological advances in neuroscience have demonstrated how information in the brain is encoded with incredibly high dimensionality with respect to both space and time (Haxby, Connolly, & Guntupalli, 2014). We believe the use of computational models will likewise better enable researchers to capture this complexity within psychological theories. Most psychological researchers are already familiar with regression as an instance of a *statistical model*, specifically a linear one that combines features according to a set of weights, in order to predict the value of a dependent variable. We encourage researchers to think of models in more general terms: a general mathematical function that transforms inputs into specific outputs². In this way, models can be likened to cooking recipes that describe how to

¹ With the clever manipulation of visual properties (e.g. color, texture, animation) it may be possible to graphically represent multiple interacting factors, but as the number of factors increase, difficulties in interpretation increase dramatically.

² This definition captures a variety of statistical and machine learning approaches including linear and nonlinear supervised, unsupervised and reinforcement learning models

best combine ingredients into prepared foods. We expand on this analogy to explicate the reasons we believe models are so critical to overcoming the Flatland fallacy.

Imagine eating a piece of cake: observing the colors and designs that draw your eye, tasting the different flavors and textures as you take a bite, and experiencing the way numerous ingredients come together to create a delightful sensory experience. What ingredients determine the colors you see and the flavors you taste? Was that a hint of vanilla? Does the frosting contain cinnamon? You might attempt to recreate this sensory experience in your own kitchen, combining numerous ingredients in different proportions, transforming those ingredients through baking at different temperatures, until you can reliably reproduce a tasty baked good. Through a painstaking trial and error process you might converge on a *recipe* for *how* to combine and transform a set of raw ingredients to produce a finished good. In much the same way that recipes serve as a set of instructions for combining specific ingredients together that produce a cake, models provide mathematical formulations for combining different input *features* together that produce an outcome. These features can span a space that comprises only a few dimensions just as it takes only takes a handful of ingredients to make a pound cake. They can also, however, be incredibly enumerate and involve complex nonlinear interactions, just as a dobos torte requires intricately interleaving many layers of a thin sponge cake with buttercream. Indeed, with recent advances in deep-learning, some models can comprise millions of distinct features organized into hidden layers that learn many different ways to filter and transform input data (e.g (Huang, Sun, Liu, Sedra, & Weinberger, 2016)). Yet despite this wide range of complexity, researchers need not manipulate so many inputs at once single handedly, instead they can rely on a model which serves as powerful and reliable assistive tool.

At their core, recipes are comprised of ingredients, proportions of ingredients, and instructions for how to combine ingredients to produce a finished product. Like recipes, models are comprised of *features* that are scaled by *weights* and combined in a specific *formula* to produce a *prediction* (Table 1). Central to our argument is that models serve as tools to both reason and communicate about *high dimensional* spaces. Models allow researchers to consider what dimensions of a problem are most relevant and predict outcomes based on complex sets of interactions. Models also allow researchers to build intuitions about their components through simulation and application to novel data sets (Yarkoni & Westfall, 2017), akin to children taking devices apart to figure out how they work. Moreover, models can be shared between researchers permitting the collective development of a cumulative science whereby weak or redundant theories are pruned and robust, predictive theories are retained.

Table 1. Relationship between Models and Recipes

<i>Recipe</i>	<i>Model</i>	<i>Purpose</i>
Ingredients	Inputs/Features	Characterize the possible building blocks (dimension) necessary to create an output
Proportions	Weights/parameters	Characterize the relative

		importance of each input
Instructions	Formula/Model specification	A describe a set of mathematical rules for ways in which inputs should be combined
Product	Prediction	The output consequent of combining inputs according to a set of rules with a set of fixed levels of importance

A key property of formalizing psychological theories as computational models is that it enables researchers to share their ideas in *extensible* ways. In contrast, the significance result of an ANOVA is fundamentally useless to other researchers trying to extend prior theoretical work. At best, researchers will only be able to devise novel experiments to test contextual moderators on the coarse treatment effects that a theory predicts (Van Bavel, Mende-Siedlecki, Brady, & Reiner, 2016). Because models are simply recipes for combining inputs to generate predictions, they can be easily applied to novel contexts using new data. For example, we have developed models of how emotions (Chang & Smith, 2015; Chang et al., 2011), social norms (Chang & Sanfey, 2013; Sanfey, Stallen, & Chang, 2014), and inferences about others (Chang, Doll, van't Wout, Frank, & Sanfey, 2010; Fareri, Chang, & Delgado, 2015; Sul, Güroğlu, Crone, & Chang, 2017) can predict decisions to cooperate. In addition, we and others have developed models for how activity in different brain regions might be combined to produce an affective response (Chang, Gianaros, Manuck, Krishnan, & Wager, 2015; Eisenbarth, Chang, & Wager, 2016; Krishnan et al., 2016; Wager et al., 2013). The advantage of these types of quantitative models is that they permit precise evaluation of how *sensitive* a model is for capturing a psychological construct, as well as how *specific* a model is to a given construct as compared to other psychological states and processes. This process of iterative construct validation through model sharing and testing on many types of data sets, is critical for developing a cumulative science of what comprises psychological states and how they are encoded and represented in the brain (Woo, Chang, Lindquist, & Wager, 2017).

Improved Quantitative Training

Although computational models offer one tool for uncovering psychological phenomena in higher dimensions, they require a dramatic reform in the way the discipline of psychology carries out quantitative training. Instead of providing a limited introduction to inference using statistics and research design and separately prioritizing the memorization of psychological effects in different domains of psychology, we believe there should be increased emphasis on teaching technical skills and a better education of data-driven inferential reasoning within all psychology courses. Moreover, psychology's curriculum could be expanded to adapt to the recent technological advances that have resulted in the exponential growth in the collection of data via the internet, online commerce, mobile sensing, etc (Griffiths, 2015; Lazer et al., 2009; Yarkoni, 2012). Beyond the narrow domain of academic psychology, there already exists

intense demand for skilled workers that can gain insights about human behavior from data in almost every industry including government, journalism, business, and healthcare. We believe that with improved training, psychologists could become increasingly involved in these efforts, ultimately providing an opportunity to inform many different and important industries.

To do so however, will require reimagining training in psychology. Working with large, complicated datasets requires basic training in areas traditionally associated with computer science and informatics including: programming, algorithms, databases, and computing. This requires extending basic education of statistical training to include skills such as data manipulation, generating predictive models, machine-learning, natural language processing, graph theory, and visualization (Montag, Duke, & Markowitz, 2016; Yarkoni, 2012; Yarkoni & Westfall, 2017). These types of competencies comprise an emerging growth of applied statistics or “data science” programs (P. Anderson, Bowring, McCauley, Pothering, & Starr, 2014). We believe that increased emphasis on training basic technical and quantitative skills will improve the ability of psychology majors to participate in the enormous endeavor of understanding human behavior from data. However, we are certainly not advocating that psychology majors should additionally pursue an accompanying degree in statistics or computer science. Instead, we recommend that training programs in psychology consider adding additional requirements to curricula (e.g. programming and computing for psychologists & advanced statistics), provide better integration of making inferences with data training into existing curricula (e.g. requiring data analysis projects), and offering more courses on advanced research methods (e.g. natural language processing, mobile sensing, or social network analysis). One practical recommendation akin to training in other STEM disciplines (e.g., biology, chemistry, physics) would be to add accompanying laboratories to the core psychology topic classes and provide hands on training for making inferences using these types of methods. We believe that this is an exciting opportunity to advance our field to new dimensions.

Towards Computational Thinking

We hesitate to leave interested readers with the intuition that improved quantitative skills and the application of computational modeling are simply additional “tools” that researchers should strive to acquire in service of conducting high dimensional psychological research. Rather, through the act of engaging in *computational thinking*, psychological researchers themselves can fundamentally change the way they approach psychological problems (N. D. Anderson, 2016).

For example, in addition to the psychological explanations enumerated above, there is also a statistical explanation of the Flatland fallacy. It is well established in the field of machine-learning that making predictions entails a trade off between minimizing bias and variance (Hastie, Tibshirani, & Friedman, 2009). Specifically, errors in predictions can be decomposed into three separate components: (1) *irreducible error* is the level of noise intrinsic to the problem, (2) *bias error* describes errors between the true data generating model and the learned model averaged across many data samples, and (3) *variance error* reflects the model’s sensitivity to the idiosyncrasies present within individual data samples (Geman, Bienenstock, & Doursat, 1992). It has been demonstrated that in the context of small sample sizes and high-dimensional signals, lower dimensional models associated with greater bias error can

actually make *more* accurate out-of-sample predictions than the true high dimensional model of an underlying signal (Friedman, 1997). This means that there can be a computational benefit to prioritizing parsimony and bias when predicting complex psychological phenomenon from small datasets. This offers an additional explanation for why so many researchers have converged on low dimensional accounts of psychological phenomena. While psychologists might be naturally inclined to further simplify their experimental manipulation and increase their sample size to improve power, computational thinking predicts adopting an alternative strategy. Rather than reducing high bias error via collecting larger sample sizes alone, computational thinking highlights the importance of measuring psychological phenomena with greater sampling *diversity*. By measuring and eliciting psychological phenomena in numerous ways, researchers can more richly sample high dimensional effects of interest. It is with this data diversity that high dimensional models can outperform biased lower dimensional alternatives. To this end, to combat the Flatland fallacy, we believe that psychological scientists should additionally strive for large sample sizes as well as large data *diversity*, richly sampling from a larger gamut of human experience.

Summary

In this paper, we have outlined how subjective biases, limitations of human cognition, social and cultural norms surrounding experimental design, analysis and pedagogy, and communicative shortcomings in discussing complex ideas can lead psychological researchers to consistently converge on low dimensional explanations of human psychology. To avoid committing this “Flatland fallacy”, we have proposed a reimagining of the field of psychology, which emphasizes a culture of developing, testing, and sharing computational models, accompanied by improved quantitative and technical training. Like cooking recipes, models provide a formal framework for taking inputs and transforming them into products (predictions). More specifically, models offer researchers a tool to assist reasoning in higher dimensional ways, approaching psychological science as an explanatory and predictive discipline (Yarkoni & Westfall, 2017), and most importantly, facilitating the development of a cumulative science rather than one characterized by a patchwork of disparate low-dimensional theories.

Acknowledgements: The authors thanks Emma Templeton, Daisy Burr, and Kristina Rapuano for feedback on earlier drafts of this manuscript. The authors also thank Seth Frey and Tal Yarkoni for early discussions of this manuscript.

Author contributions: E.J. and L.J.C conceived of and wrote the manuscript.

Competing financial interests: The authors declare no competing financial interests.

References

- Abbott, E. A. (1884). *Flatland: A Romance of Many Dimensions*, by A. Square. United Kingdom: Seeley & Co.
- Anderson, N. D. (2016). A Call for Computational Thinking in Undergraduate Psychology. *Psychology Learning & Teaching*, *15*(3), 226–234.
- Anderson, P., Bowring, J., McCauley, R., Pothering, G., & Starr, C. (2014). An Undergraduate Degree in Data Science: Curriculum and a Decade of Implementation Experience. In *Proceedings of the 45th ACM Technical Symposium on Computer Science Education* (pp. 145–150). New York, NY, USA: ACM.
- Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review*, *105*(3), 442–481.
- Bolles, R. C. (1962). The difference between statistical hypotheses and scientific hypotheses. *Pharmacological Reports: PR*, *11*(7), 639.
- Brunswick, E. (1952). *The Conceptual Framework of Psychology*. Chicago; London; Toronto: University of Chicago Press.
- Camerer, C. F., Ho, T.-H., & Chong, J.-K. (2004). A Cognitive Hierarchy Model of Games. *The Quarterly Journal of Economics*, *119*(3), 861–898.
- Camerer, C. F., Ho, T.-H., & Chong, J. K. (2015). A psychological approach to strategic thinking in games. *Current Opinion in Behavioral Sciences*, *3*, 157–162.
- Camerer, C. F., Johnson, E., Rymon, T., & Sen, S. (1993). Cognition and framing in sequential bargaining for gains and losses. *Frontiers of Game Theory*, *1*, 27–47.
- Chaiken, S., & Trope, Y. (1999). *Dual-process Theories in Social Psychology*. New York, NY: Guilford Press.
- Chang, L. J., Doll, B. B., van't Wout, M., Frank, M. J., & Sanfey, A. G. (2010). Seeing is believing: Trustworthiness as a dynamic belief. *Cognitive Psychology*, *61*(2), 87–105.
- Chang, L. J., Gianaros, P. J., Manuck, S. B., Krishnan, A., & Wager, T. D. (2015). A Sensitive and Specific Neural Signature for Picture-Induced Negative Affect. *PLoS Biology*, *13*(6), e1002180.
- Chang, L. J., & Sanfey, A. G. (2008). Emotion, decision-making and the brain. In D. Houser & K. McCabe (Eds.), *Neuroeconomics* (Vol. 20, pp. 31–53). United Kingdom.
- Chang, L. J., & Sanfey, A. G. (2013). Great expectations: neural computations underlying the use of social norms in decision-making. *Social Cognitive and Affective Neuroscience*, *8*(3), 277–284.
- Chang, L. J., & Smith, A. (2015). Social emotions and psychological games. *Current Opinion in Behavioral Sciences*, *5*, 133–140.
- Chang, L. J., Smith, A., Dufwenberg, M., & Sanfey, A. G. (2011). Triangulating the neural, psychological, and economic bases of guilt aversion. *Neuron*, *70*(3), 560–572.
- Cohen, A. R., Stotland, E., & Wolfe, D. M. (1955). An experimental investigation of need for cognition. *Journal of Abnormal Psychology*, *51*(2), 291–294.
- Costa-Gomes, M., Crawford, V. P., & Broseta, B. (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica: Journal of the Econometric Society*, *69*(5), 1193–1235.
- Cuddy, A. J. C., Fiske, S. T., & Glick, P. (2008). Warmth and Competence as Universal Dimensions of Social Perception: The Stereotype Content Model and the BIAS Map. *Advances in Experimental Social Psychology*, *40*, 61–149.
- Damasio, A. R. (1994). *Descartes' Error: Emotion, Rationality and the Human Brain*. New York:

NY: Harper Collins.

- Davidson, R. J. (1993). Cerebral asymmetry and emotion: Conceptual and methodological conundrums. *Cognition and Emotion*, *7*(1), 115–138.
- Dawes, R. M. (1971). A case study of graduate admissions: Application of three principles of human decision making. *The American Psychologist*, *26*(2), 180–188.
- Dawes, R. M., Faust, D., & Meehl, P. E. (1989). Clinical versus actuarial judgment. *Science*, *243*(4899), 1668–1674.
- Daw, N. D., Niv, Y., & Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nature Neuroscience*, *8*(12), 1704–1711.
- Deane, D. H., Hammond, K. R., & Summers, D. A. (1972). Acquisition and application of knowledge in complex inference tasks. *Journal of Experimental Psychology*, *92*(1), 20–26.
- Dijksterhuis, A., Bos, M. W., Nordgren, L. F., & van Baaren, R. B. (2006). On making the right choice: the deliberation-without-attention effect. *Science*, *311*(5763), 1005–1007.
- Dudycha, L. W., & Naylor, J. C. (1966). Characteristics of the human inference process in complex choice behavior situations. *Organizational Behavior and Human Performance*, *1*(1), 110–128.
- Eisenbarth, H., Chang, L. J., & Wager, T. D. (2016). Multivariate Brain Prediction of Heart Rate and Skin Conductance Responses to Social Threat. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *36*(47), 11987–11998.
- Epstein, S. (1994). Integration of the cognitive and the psychodynamic unconscious. *The American Psychologist*, *49*(8), 709–724.
- Evans, J. (2003). In two minds: dual-process accounts of reasoning. *Trends in Cognitive Sciences*, *7*(10), 454–459.
- Evans, J., & Frankish, K. (2009). *In Two Minds: Dual Processes and Beyond*. Oxford, United Kingdom: Oxford University Press.
- Fareri, D. S., Chang, L. J., & Delgado, M. R. (2015). Computational substrates of social value in interpersonal collaboration. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *35*(21), 8170–8180.
- Fisher, R. A. (1925). Statistical Methods for Research Workers. In *Breakthroughs in Statistics* (pp. 66–70). Springer, New York, NY.
- Frank, M. J., Cohen, M. X., & Sanfey, A. G. (2009). Multiple Systems in Decision Making. *Current Directions in Psychological Science*, *18*(2), 73–77.
- Friedman, J. H. (1997). On bias, variance, 0/1—loss, and the curse-of-dimensionality. *Data Mining and Knowledge Discovery*.
- Geman, S., Bienenstock, E., & Doursat, R. (1992). Neural networks and the bias/variance dilemma. *Neural Computation*, *4*, 1–58.
- Gigerenzer, G., & Brighton, H. (2009). Homo heuristics: why biased minds make better inferences. *Topics in Cognitive Science*, *1*(1), 107–143.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: models of bounded rationality. *Psychological Review*, *103*(4), 650–669.
- Gilbert, D. T., & Wilson, T. D. (2000). Miswanting: Some Problems in the Forecasting of Future Affective States. In J. P. Forgas (Ed.), *Feeling and Thinking: The role of affect in social cognition* (Vols. 1–8, pp. 178–197). Cambridge: Cambridge Univ Press.
- Gray, H. M., Gray, K., & Wegner, D. M. (2007). Dimensions of mind perception. *Science*, *315*(5812), 619.
- Gray, K. (2017). How to Map Theory: Reliable Methods Are Fruitless Without Rigorous Theory. *Perspectives on Psychological Science: A Journal of the Association for Psychological*

- Science*, 1745691617691949.
- Gray, K., & Wegner, D. M. (2009). Moral typecasting: divergent perceptions of moral agents and moral patients. *Journal of Personality and Social Psychology*, 96(3), 505–520.
- Greene, J. D., Nystrom, L. E., Engell, A. D., Darley, J. M., & Cohen, J. D. (2004). The neural bases of cognitive conflict and control in moral judgment. *Neuron*, 44(2), 389–400.
- Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M., & Cohen, J. D. (2001). An fMRI investigation of emotional engagement in moral judgment. *Science*, 293(5537), 2105–2108.
- Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: attitudes, self-esteem, and stereotypes. *Psychological Review*, 102(1), 4–27.
- Griessinger, T., & Coricelli, G. (2015). The neuroeconomics of strategic interaction. *Current Opinion in Behavioral Sciences*, 3(Supplement C), 73–79.
- Griffiths, T. L. (2015). Manifesto for a new (computational) cognitive revolution. *Cognition*, 135, 21–23.
- Happé, F. G. (1994). An advanced test of theory of mind: understanding of story characters' thoughts and feelings by able autistic, mentally handicapped, and normal children and adults. *Journal of Autism and Developmental Disorders*, 24(2), 129–154.
- Haslam, N. (2006). Dehumanization: an integrative review. *Personality and Social Psychology Review: An Official Journal of the Society for Personality and Social Psychology, Inc.*, 10(3), 252–264.
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The Elements of Statistical Learning; Data Mining, Inference and Prediction*. New York, NY: Springer.
- Haxby, J. V., Connolly, A. C., & Guntupalli, J. S. (2014). Decoding neural representational spaces using multivariate pattern analysis. *Annual Review of Neuroscience*, 37, 435–456.
- Heatherton, T. F., & Wagner, D. D. (2011). Cognitive neuroscience of self-regulation failure. *Trends in Cognitive Sciences*, 15(3), 132–139.
- Heusser, A. C., Ziman, K., Owen, L. L. W., & Manning, J. R. (2017, January 28). *HyperTools: A Python toolbox for visualizing and manipulating high-dimensional data*. *arXiv [stat.OT]*. Retrieved from <http://arxiv.org/abs/1701.08290>
- Hofmann, W., Friese, M., & Strack, F. (2009). Impulse and Self-Control From a Dual-Systems Perspective. *Perspectives on Psychological Science: A Journal of the Association for Psychological Science*, 4(2), 162–176.
- Huang, G., Sun, Y., Liu, Z., Sedra, D., & Weinberger, K. (2016, March 30). *Deep Networks with Stochastic Depth*. *arXiv [cs.LG]*. Retrieved from <http://arxiv.org/abs/1603.09382>
- Kahneman, D. (2003). A perspective on judgment and choice: mapping bounded rationality. *The American Psychologist*, 58(9), 697–720.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. New York, NY: Farrar, Straus and Giroux.
- Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica: Journal of the Econometric Society*, 47(2), 263–291.
- Karelaia, N., & Hogarth, R. M. (2008). Determinants of linear judgment: a meta-analysis of lens model studies. *Psychological Bulletin*, 134(3), 404–426.
- Karlsson, F. (2007). Constraints on multiple center-embedding of clauses. *Journal of Linguistics*, 43(02), 365.
- Karlsson, F. (2010). Syntactic recursion and iteration. In H. van der Hulst & H. van der Hulst (Eds.), *Recursion and Human Language*. Berlin, New York: De Gruyter Mouton.
- Kerr, N. L. (1998). HARKing: hypothesizing after the results are known. *Pers. Soc. Psychol. Rev.*, 2(3), 196–217.
- Koriat, A. (1993). How do we know that we know? The accessibility model of the feeling of knowing. *Psychological Review*, 100(4), 609–639.

- Krishnan, A., Woo, C.-W., Chang, L. J., Ruzic, L., Gu, X., López-Solà, M., ... Wager, T. D. (2016). Somatic and vicarious pain are represented by dissociable multivariate brain patterns. *eLife*, 5. <https://doi.org/10.7554/eLife.15166>
- Lazer, D., Pentland, A., Adamic, L., Aral, S., Barabasi, A.-L., Brewer, D., ... Van Alstyne, M. (2009). Social science. Computational social science. *Science*, 323(5915), 721–723.
- McClure, S. M., Laibson, D. I., Loewenstein, G., & Cohen, J. D. (2004). Separate neural systems value immediate and delayed monetary rewards. *Science*, 306(5695), 503–507.
- Medin, D. L. (1989). Concepts and conceptual structure. *The American Psychologist*, 44(12), 1469–1481.
- Meehl, P. E. (1954). *Clinical versus statistical prediction: A theoretical analysis and a review of the evidence*. Minneapolis, MN, US: University of Minnesota Press.
- Meehl, P. E. (1967). Theory-Testing in Psychology and Physics: A Methodological Paradox. *Philosophy of Science*, 34(2), 103–115.
- Meehl, P. E. (1990). Why summaries of research on psychological theories are often uninterpretable. *Psychological Reports*, 66(1), 195.
- Mervis, C. B., & Rosch, E. (1981). Categorization of Natural Objects. *Annual Review of Psychology*, 32(1), 89–115.
- Miller, G. A. (1956). The magic number seven plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63, 91–97.
- Mitchell, J. P. (2005). The false dichotomy between simulation and theory-theory: the argument's error. *Trends in Cognitive Sciences*, 9(8), 363–4; author reply 364.
- Montag, C., Duke, É., & Markowitz, A. (2016). Toward Psychoinformatics: Computer Science Meets Psychology. *Computational and Mathematical Methods in Medicine*, 2016, 2983685.
- Montague, R. (2007). *Your Brain Is (Almost) Perfect: How We Make Decisions*. New York, NY: Penguin.
- Ochsner, K. N., Bunge, S. A., Gross, J. J., & Gabrieli, J. D. E. (2002). Rethinking feelings: an fMRI study of the cognitive regulation of emotion. *Journal of Cognitive Neuroscience*, 14(8), 1215–1229.
- Pashler, H. (1994). Dual-task interference in simple tasks: data and theory. *Psychological Bulletin*, 116(2), 220–244.
- Payne, J. W. (1976). Task complexity and contingent processing in decision making: An information search and protocol analysis. *Organizational Behavior and Human Performance*, 16(2), 366–387.
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. New York, NY: Cambridge University Press.
- Pinker, S. (1999). How the mind works. *Annals of the New York Academy of Sciences*, 882, 119–27; discussion 128–34.
- Poldrack, R. A., Clark, J., Paré-Blagoev, E. J., Shohamy, D., Creso Moyano, J., Myers, C., & Gluck, M. A. (2001). Interactive memory systems in the human brain. *Nature*, 414(6863), 546–550.
- Poldrack, R. A., & Packard, M. G. (2003). Competition among multiple memory systems: converging evidence from animal and human brain studies. *Neuropsychologia*, 41(3), 245–251.
- Posner, J., Russell, J. A., & Peterson, B. S. (2005). The circumplex model of affect: an integrative approach to affective neuroscience, cognitive development, and psychopathology. *Development and Psychopathology*, 17(3), 715–734.
- Rand, D. G., Greene, J. D., & Nowak, M. A. (2012). Spontaneous giving and calculated greed. *Nature*, 489(7416), 427–430.

- Rozenblit, L., & Keil, F. (2002). The misunderstood limits of folk science: an illusion of explanatory depth. *Cognitive Science*, 26(5), 521–562.
- Russell, J. A. (1980). A circumplex model of affect. *Journal of Personality and Social Psychology*, 39(6), 1161–1178.
- Sanfey, A. G., & Chang, L. J. (2008). Multiple systems in decision making. *Annals of the New York Academy of Sciences*, 1128, 53–62.
- Sanfey, A. G., Stallen, M., & Chang, L. J. (2014). Norms and expectations in social decision-making. *Trends in Cognitive Sciences*, 18(4), 172–174.
- Saxe, R. (2005). Against simulation: the argument from error. *Trends in Cognitive Sciences*, 9(4), 174–179.
- Schachter, S., & Singer, J. E. (1962). Cognitive, social, and physiological determinants of emotional state. *Psychological Review*, 69, 379–399.
- Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review*, 84(1), 1–66.
- Shiffrin, R. M., & Schneider, W. (1984). Automatic and controlled processing revisited. *Psychological Review*, 91(2), 269–276.
- Simon, H. A. (1957). *Models of Man: Social and Rational*. New York, NY: Wiley.
- Simon, H. A. (1987). Satisficing. In P. Macmillan (Ed.), *The New Palgrave Dictionary of Economics* (pp. 1–3). Palgrave Macmillan UK.
- Slooman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, 119(1), 3–22.
- Smith, E. R., & Zarate, M. A. (1992). Exemplar-based model of social judgment. *Psychological Review*, 99(1), 3.
- Stahl, D. O., & Wilson, P. W. (1995). On Players' Models of Other Players: Theory and Experimental Evidence. *Games and Economic Behavior*, 10(1), 218–254.
- Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: implications for the rationality debate? *The Behavioral and Brain Sciences*, 23(5), 645–65; discussion 665–726.
- Steele, C. M., & Aronson, J. (1995). Stereotype threat and the intellectual test performance of African Americans. *Journal of Personality and Social Psychology*, 69(5), 797–811.
- Stoloff, M., McCarthy, M., Keller, L., Varfolomeeva, V., Lynch, J., Makara, K., ... Smiley, W. (2009). The Undergraduate Psychology Major: An Examination of Structure and Sequence. *Teaching of Psychology*, 37(1), 4–15.
- Stevens, M. (2000). The essentialist aspect of naive theories. *Cognition*, 74(2), 149–175.
- Sul, S., Güroğlu, B., Crone, E. A., & Chang, L. J. (2017). Medial prefrontal cortical thinning mediates shifts in other-regarding preferences during adolescence. *Scientific Reports*, 7(1), 8510.
- Timmermans, D. (1993). The impact of task complexity on information use in multi-attribute decision making. *Journal of Behavioral Decision Making*, 6(2), 95–111.
- Todorov, A., Said, C. P., Engell, A. D., & Oosterhof, N. N. (2008). Understanding evaluation of faces on social dimensions. *Trends in Cognitive Sciences*, 12(12), 455–460.
- Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, 185(4157), 1124–1131.
- Van Bavel, J. J., Mende-Siedlecki, P., Brady, W. J., & Reinero, D. A. (2016). Contextual sensitivity in scientific reproducibility. *Proceedings of the National Academy of Sciences of the United States of America*, 113(23), 6454–6459.
- van der Maaten, L., & Hinton, G. (2008). Visualizing Data using t-SNE. *Journal of Machine Learning Research: JMLR*, 9, 2579–2605.
- Van Horn, J. D., & Toga, A. W. (2014). Human neuroimaging as a “Big Data” science. *Brain*

- Imaging and Behavior*, 8(2), 323–331.
- Wager, T. D., Atlas, L. Y., Lindquist, M. A., Roy, M., Woo, C.-W., & Kross, E. (2013). An fMRI-based neurologic signature of physical pain. *The New England Journal of Medicine*, 368(15), 1388–1397.
- Watts, D. J. (2017). Should social science be more solution-oriented? *Nature Human Behaviour*, 1(1), s41562–016–0015.
- Waytz, A., & Mitchell, J. P. (2011). Two mechanisms for simulating other minds: dissociations between mirroring and self-projection. *Current Directions in Psychological Science*, 20(3), 197–200.
- Webster, D. M., & Kruglanski, A. W. (1994). Individual differences in need for cognitive closure. *Journal of Personality and Social Psychology*, 67(6), 1049–1062.
- Wilson, T. D., & Schooler, J. W. (1991). Thinking too much: introspection can reduce the quality of preferences and decisions. *Journal of Personality and Social Psychology*, 60(2), 181–192.
- Woo, C.-W., Chang, L. J., Lindquist, M. A., & Wager, T. D. (2017). Building better biomarkers: brain models in translational neuroimaging. *Nature Neuroscience*, 20(3), 365–377.
- Yarkoni, T. (2012). Psychoinformatics: New horizons at the interface of the psychological and computing sciences. *Current Directions in Psychological Science*, 21(6), 391–397.
- 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*, 1745691617693393.
- Zajonc, R. B. (1980). Feeling and thinking: Preferences need no inferences. *The American Psychologist*, 35(2), 151–175.