American Psychologist

ISSN: 0003-066X

2021, Vol. 76, No. 6, 1013–1026
https://doi.org/10.1037/amp0000847

The Cultural Dynamics of Concept Creep

Nick Haslam[1], Ekaterina Vylomova[1], Michael Zyphur[2], and Yoshihisa Kashima[1]
[1] Melbourne School of Psychological Sciences, University of Melbourne
[2] Faculty of Business and Economics, University of Melbourne

Emerging methods for studying cultural dynamics allow researchers to investigate cultural change with newfound rigor. One change that has recently attracted the attention of social commentators is "concept creep," the semantic inflation of harm-related concepts such as trauma, bullying, and prejudice. In theory, concept creep is driven distally by several recent cultural and societal trends, but psychology also plays a proximal role in developing and disseminating expansionary concepts of harm. However, there have been few systematic attempts to document concept creep and none to explore factors that influence it. The present work reviews concept creep from the perspective of cultural dynamics and lays out a conceptual framework for exploring processes implicated in it. Illustrative analyses are presented that apply computational linguistic methods to very large text corpora, including a new corpus of psychology article abstracts. They demonstrate that harm has risen steeply in prominence both in psychology and in the wider culture in recent decades, and that harm-related concepts have inflated their meanings over this period. The analyses also provide evidence of dynamic relationships between the prominence and semantic breadth of harm-related concepts, and between psychology and the culture at large. Implications are drawn for theory and research on concept creep.

Public Significance Statement
This study indicates that concepts of harm have broadened their meanings and become more prominent in psychology over the past half century. It suggests that similar changes have occurred in the culture at large, and that the respective changes may be dynamically linked. These findings signal that psychology is implicated in an important cultural shift.

Keywords: concept creep, culture, dynamics, harm, morality

Supplemental materials: https://doi.org/10.1037/amp0000847.supp

Culture is increasingly conceptualized as a distribution of socially transmissible information. This conceptualization is now prominent in psychology (e.g., Chiu & Hong, 2013; Heine, 2020; Kashima, 2008), anthropology (e.g., Boyd & Richerson, 1985; Sperber, 1996), and biology (e.g., Cavalli-Sforza & Feldman, 1981; Mesoudi, 2011). Understood in this fashion, cultural dynamics are the processes by which information is created or imported, transmitted or remembered, and

Editor's note. This article is part of the special issue "Psychological Perspectives on Cultural Change" published in the September 2021 issue of *American Psychologist*. Michael E. W. Varnum and Igor Grossmann served as guest editors of the special issue, with Lillian Comas-Díaz as advisory editor.

Nick Haslam https://orcid.org/0000-0002-1913-2340
Yoshihisa Kashima https://orcid.org/0000-0003-3627-3273

Nick Haslam served as lead for data curation and software, contributed equally to formal analysis and validation, and served in a supporting role for conceptualization, methodology, administration, and writing–review and editing.

Ekaterina Vylomova served as lead for data curation and software, contributed equally to validation and served in a supporting role for writing–review and editing. Michael Zyphur contributed equally to validation and served in a supporting role for writing–review and editing. Yoshihisa Kashima contributed equally to validation and served in a supporting role for writing–review and editing.

The research reported in this article was supported by Australian Research Council Discovery Projects DP170104948 and DP210103984.

Correspondence concerning this article should be addressed to Nick Haslam, Melbourne School of Psychological Sciences, University of Melbourne, Parkville, VIC 3010, Australia. Email: nhaslam@unimelb.edu.au

transformed or forgotten within a society. These processes unfold in many ways and at many levels, from the interpersonal communication of individuals to the tectonic shifts of whole societies (Kashima et al., 2019). If the study of culture is, as Sperber (1985) argued, the "epidemiology of representations" (Sperber, 1985), then cultural change is akin to the spread of viruses (e.g., Lamberson, 2016) or the diffusion of technological innovations (e.g., Rogers, 2003). Cultural epidemiologists can investigate what information spreads, how fast it spreads, and who spreads it.

The psychology of cultural dynamics has itself spread rapidly (for reviews, see Kashima et al., 2019; Varnum & Grossmann, 2017), in part aided by new conceptual tools and data-intensive diachronic methodologies. Although some have studied cultural persistence (e.g., Alesina et al., 2013), many researchers have documented and sought to explain shifts on broad cultural dimensions such as individualism (e.g., Greenfield, 2013; Hamamura, 2012), as well as in the meaning of more specific concepts such as happiness (Oishi et al., 2013) and nature (e.g., Kesebir & Kesebir, 2017; Wolff et al., 1999). They have plotted transformations at timescales of decades or centuries. These shifts are assessed as historical changes in the prevalence or salience of classes of ideas or practices indexed from archival or behavioral trace data (Kashima, 2014). Common examples include the relative frequency of words or phrases of interest (e.g., baby-names) in large text corpora (e.g., the "culturomics" approach launched by Michel et al., 2011), or average scores on self-report measures of psychological properties.

For example, rising individualism—one of the focal points of investigation—might be inferred from increasing frequencies of first-person pronouns (e.g., Hamamura & Xu, 2015; Twenge et al., 2012) or words associated with individualism in the Google Books corpus (e.g., Greenfield, 2013; Grossmann & Varnum, 2015); from cohort effects in scores on narcissism scales in a cross-temporal meta-analysis (Twenge & Foster, 2010) or from shifting life aspirations and self-views taken from yearly surveys of American high school students and undergraduates (e.g., Twenge & Campbell, 2008; Twenge et al., 2012). These methods of quantifying alterations in the prevalence, salience, or endorsement of ideas and practices allow the trajectories of cultural changes to be described and their potential determinants to be evaluated statistically. In the case of individualism, recent research has highlighted the role of socioecological factors such as economic wealth and urbanization, disaster and pathogen prevalence, and the spread of marriage and family practices encouraged by the Catholic church as potential drivers (e.g., Grossmann & Varnum, 2015; Santos et al., 2017; Schulz et al., 2019; Varnum & Grossmann, 2017).

This dynamic approach may shine a light on "concept creep" (Haslam, 2016; Haslam et al., 2020). This term identifies a tendency for concepts related to harm—that is, those associated with suffering, damage, and destruction on the one hand, and associated concepts of care, safety, and protection on the other (cf. Moral Foundations Theory; Graham et al., 2009)—to have broadened their meanings in recent decades. For example, concepts such as trauma and bullying have crept if they now refer to a wider range of phenomena than they did in earlier times (e.g., less severe or vicarious experiences of adversity coming to be conceptualized as traumas or online intimidation being added to the definition of bullying). Concept creep has received significant attention from social commentators in recent years in connection with the so-called culture wars, recent upheavals on college campuses (e.g., Campbell & Manning, 2018; Lukianoff & Haidt, 2018), and debates over contentious ideas such as trigger warnings, safe spaces, microaggressions, and the limits of free speech. However, it has received vastly less attention from researchers than more well-established cultural dimensions. There has been almost no work to date investigating concept creep as a diachronic phenomenon, despite it being framed explicitly as a historical change process. The emerging psychology of cultural dynamics should provide a toolkit for making sense of this process. This article presents an overview of theory and research on concept creep, explores how the cultural dynamics approach might illuminate it, and provides a preliminary, illustrative analysis that demonstrates how that approach might be addressed in a rigorous research program, using the tools of computational linguistics and time series analysis.

Applying a cultural dynamics lens to concept creep may not only refine our understanding of the phenomenon but also contribute to the elaboration of how we understand cultural dynamics. Compared with most recent psychological work in this field, research on concept creep has a few points of distinctiveness. First, concept creep is a relatively circumscribed cultural-historical phenomenon, involving a specific set of concepts changing in a particular way over a few decades. In this respect it contrasts with recent psychological research on changes in broad cultural dimensions such as individualism-collectivism (Kashima, 2001). Second, concept creep directly implicates changes in semantic representations, whereas changes in individualism and related cultural dimensions primarily refer to values and practices rather than word meanings (although it is possible that changes in individualism may be reflected in changes in the meaning of the individual). It is particularly well suited to an "epidemiology of representations" approach. Third, whereas most quantitative studies of cultural dynamics assess them in terms of the prevalence of particular content, as in the word frequency-based methods used in "culturomics" (Michel et al., 2011), concept creep primarily involves changes in the meaning rather than prevalence or prominence of ideas. Another way to express this point is

that past cultural dynamics research has generally focused on people's changing psychological relations to particular referents (i.e., the people, things, events, and situations to which words or other signs refer) and assumed that the sign-referent relations are stable over time. In contrast, concept creep research is concerned with historical shifts in people's understandings of sign-referent relations: how words can come to signify a broadening range of referents. Finally, the theory of concept creep is distinctive in the context of prior research on cultural dynamics in proposing somewhat novel influences on the phenomenon. Whereas accounts of changes in some broad cultural dimensions may invoke socioecological determinants, concept creep is attributed in part to the diffusion of ideas from the academic discourse of psychology into the wider culture. Although concept creep might originate from other sources—political elites, thought leaders, social media, other corners of the academy—diffusion from our discipline has particular importance to psychologists. Studying concept creep from the perspective of cultural dynamics affords the possibility of examining new kinds of dynamic influences on culture.

Concept Creep: Theory

Characterizing Concept Creep

Haslam (2016) initially laid out his account of concept creep by presenting case studies of six concepts in widespread use within psychology. The concepts were diverse, falling within the domains of developmental and organizational psychology (abuse, bullying), clinical psychology (addiction, mental disorder, trauma), and social psychology (prejudice), and were intended as illustrative examples rather than an exhaustive set. Haslam presented historical analyses of the use and definition of these concepts within psychology and cognate fields to argue that their meanings had progressively expanded over time. For example, when "bullying" was introduced to psychology in the 1970s it was applied to a specific form of peer aggression among children, restricted to behavior that was repeated, intentional, and carried out in the context of a power imbalance where the perpetrator dominated the victim in age, size, or number, for example. In later decades, bullying scholarship gradually relaxed each of these conditions: bullying could be identified in cases where behavior was unrepeated, unintentional, and carried out by a perpetrator of equal or lower power or status. The behavior recognized as bullying increasingly included acts of omission such as shunning or passive exclusion, and bullying also came to be used to refer to inappropriate behavior among adults in workplaces. As a result, the term encompassed a much broader range of acts and experiences than it had at its inception. It is this broadening of the range of referents to which harm-related concepts apply that is the essence of concept creep.

Haslam (2016) offered similar accounts of the semantic inflation of the other putatively creeping concepts. Active "abuse" increasingly came to incorporate passive neglect, and forms of "emotional abuse" were added to the primarily physical forms that were initially recognized. The sphere of mental disorders expanded with each successive edition of the *Diagnostic and Statistical Manual of Mental Disorders* (*DSM*), and criteria for some disorders became more relaxed and inclusive so more people met them. Definitions of trauma broadened the range of traumatic events recognized as possible occasions for posttraumatic stress disorder from experiences that were personally life-threatening to those that were witnessed or experienced vicariously. Addictions expanded to accommodate behavioral addictions involving compulsive gambling or shopping rather than being restricted to dependency on ingested substances. The meaning of prejudice grew to include animosities to a wider range of groups and was also expanded to include unconscious, aversive, or "modern" forms in addition to overt bigotries. The semantic enlargement revealed by these examples took two main forms, according to Haslam (2016). In "vertical creep" concept meanings became less stringent by a relaxation of criteria or a lowering of a threshold, such that less severe or intense phenomena came to be regarded as examples of the concept (e.g., vicarious trauma, unrepeated bullying). As a result, the concept comes to be used in ways that seem hyperbolic from the perspective of its earlier meaning. In "horizonal creep," by contrast, concepts broaden by incorporating qualitatively different and new phenomena (e.g., addition of new domains of psychopathology in *DSM*, recognition of new targets of prejudice). Both forms of creep can be distinguished from the emergence of polysemy—the existence of multiple distinct meanings for a single word, such as "wood" as both a substance and as a heavily treed area—because the expanded meaning remains singular but covering a wider range of referents. Equally, concept creep should be distinguished from the more general phenomenon of semantic broadening—that is not restricted to the harm-related concepts explored by Haslam (2016)—and from the even more general phenomenon of semantic change (Hamilton et al., 2016), which can involve semantic narrowing and drift as well as broadening.

Explaining Concept Creep

In addition to characterizing concept creep, Haslam (2016) proposed an account of the factors contributing to it. He argued that although the concepts examined in his six case studies were disparate, they all shared a thematic element of harm, referring to ways in which people experience harm, harm others, or are harmed by others. Seeking a parsimonious account of concept creep, rather than a set of concept-specific accounts (e.g., "medicalization" as an

explanation for the expansion of the concept of mental disorder), Haslam proposed that the semantic expansion of harm-related concepts reflected a growing cultural sensitivity to harm. In assigning a central place to harm in concept creep, Haslam was not claiming that semantic expansion is in any way specific or unique to harm-related concepts, or that all such concepts have inflated, but merely proposing a noteworthy and consequential trend and giving it a name. To be explicit, concept creep occurs when concepts associated with harm in the sense used by Moral Foundations Theory (a theory proposing five distinct moral domains; Graham et al., 2009)—that is, both negatively valenced (harm in the narrow sense of suffering, damage, and destruction) and positively valenced (care-based concepts such as empathy, safety, and protection; concepts that revolve around the prevention or palliation of harm)—expand their meanings. Haslam's proposal was largely agnostic as to the drivers of that increased sensitivity but followed Pinker's (2011) analysis of recent declines in rates of violence by invoking the rights revolutions of the 1960s. On this account, broadened meanings of abuse, bullying, and prejudice redefine the threshold of unacceptable behavior so that previously tolerated acts are problematized and penalized.

A more recent review of concept creep (Haslam et al., 2020) developed a more systematic account of its causes, proposing three primary determinants. First, concept creep might be the unintended consequence of a general trend toward greater objective material comfort in Western societies over the past half century (Pinker, 2018). As exposure to more severe and blatant forms of violence, crime, oppression, and hardship has tended to decline—not for all people or at all times, needless to say—more moderate and subtle experiences have come to be defined as harmful. A compelling experimental analogue of this dynamic was identified as "prevalence-induced concept change" by Levari and colleagues (Levari et al., 2018). As blue dots, angry faces, and unethical research proposals became scarcer in their studies, participants came to classify stimuli they had previously judged to be purple, neutral, and ethical as examples of these categories. A decrease in the prevalence of a phenomenon led to a compensatory shift in the threshold for identifying it, just as a decline in the prevalence of offensive behavior might drive an expansion of what counts as offensive.

The second determinant of concept creep proposed by Haslam et al. (2020) invoked cultural change rather than alterations in the objective prevalence of societal harms. Population-level changes in values, for example, might have the effect of sensitizing people to harms that were believed to be innocuous or unremarkable in earlier times. The documented rise of postmaterialist values in the West (Inglehart, 2008); for example, represents an increased concern with quality of life relative to material security, and might encourage a reduced tolerance for suffering. Shifts in moral values might have similar effects. Graham and colleagues (Graham et al., 2011) identified Harm/Care as one of five distinct moral foundations by which people evaluate rightness and wrongness. This foundation foregrounds the harms caused to victims of immoral action and the virtue of empathic concern and nurturing those who suffer. Any rise in the cultural prominence of this foundation in recent decades could have the effect of enlarging definitions of harm. As endorsement of the Harm/Care foundation is associated with political liberalism and female gender (Graham et al., 2009), liberalizing trends within Western societies and increased female representation within their institutions might drive such cultural shifts. An alternative account, related to this one, is that such liberalizing trends in the modern political, legal, and media landscapes have afforded the expression of a wider range of views, including broadening usages and definitions of harm-related terms.

A third and final explanation for broadening concepts of harm implicates the motivated efforts of political actors. Sunstein (2018), for example, explored concept creep as an outcome of deliberate interventions by "opprobrium entrepreneurs" who stretch the boundaries of harm-related concepts so as to condemn actions they object to and stigmatize those responsible for them. By labeling behavior as bullying or speech as hateful, even when they do not reach the threshold set by existing definitions, these expansionary classifications boost the perceived seriousness of the phenomenon, tarnish the reputations of the supposed bullies and hate-mongers, and burnish the reputations of their supposed victims. The negative moral charge associated with the existing concept (e.g., racism, abuse) becomes attached to less severe phenomena for political ends. The longer-term effects of this altered labeling may be to lower the public's thresholds for identifying harm.

The societal and cultural bases for concept creep outlined above, invoking changing socioeconomic conditions and cultural values, are similar in nature to the proposed determinants of other cultural changes, such as the possible causes of global increases in individualism (Santos et al., 2017). In contrast, explanations of concept creep that invoke strategically motivated conceptual expansions are somewhat novel, and probably more plausible as accounts of specific conceptual shifts rather than of broad cultural transformations. However, Haslam's (2016) account of concept creep also proposes another novel element that promotes harm inflation, namely the discourse of psychology. Although the distal determinants of a rising concern with harm may involve deep societal transformations, more proximally many of the concepts whose growing prevalence and expanding meanings express that rising concern come from the behavioral sciences.

The case studies discussed in Haslam (2016), for example, examined concepts that have been staples of the literature of psychology and cognate fields such as psychiatry. Psychology has arguably played an active and culturally influential role in foregrounding harm by drawing systematic attention to forms of suffering (e.g., mental disorder, addiction), adversity (e.g., trauma), and destructive behavior (e.g., abuse, bullying, or prejudice). The field has problematized these forms of harm, developed a conceptual apparatus for studying them, and pursued an applied agenda of harm reduction. The harm-related concepts developed and elaborated by psychologists have diffused into everyday language and affected how people make sense of their experience and shape their identities (Hacking, 1995). Psychology is not the only academic field to prioritize harm and to disseminate harm-related concepts into the culture at large, but as earlier critiques of "therapeutic culture" (e.g., Furedi, 2004; Lasch, 1979) and more recent (e.g., De Vos, 2013) critiques of psychologization have argued, the field has had significant cultural impacts in the West. The evolving discourse of psychology, and how it has played an influencing role in wider cultural changes, is an important aspect of the cultural dynamics of concept creep.

Concept Creep: Research

Most of the empirical research inspired by the theory of concept creep has examined synchronic phenomena analogous or loosely connected to it rather than examining concept creep itself, which is an intrinsically diachronic phenomenon. This research explores the individual difference correlates of holding broad concepts of harm. Although psychological variations that predict "harm concept breadth" (HCB) cross-sectionally may not directly enlighten us about the processes or factors that contribute to the broadening of these concepts over historical time, the diachronic focus of cultural dynamics research, they may also hold clues to those processes and factors.

Several studies have demonstrated that HCB can be measured reliably and that variations in it are associated with a variety of personality traits, attributes, and demographic characteristics. McGrath et al. (2019) and McGrath and Haslam (2020) developed scales to assess HCB by presenting people with vignettes describing marginal examples of several harm-related concepts. McGrath and Haslam (2020) showed that individual differences in the breadth of disparate concepts were underpinned by a single factor, consistent with the view that the concepts share harm as a common element. These two studies also established a replicated web of correlates of HCB. People with relatively broad or inclusive concepts of harm tended to be high in affective empathy, had a liberal political orientation, and endorsed the Harm moral foundation (Graham et al., 2009). They tended to be sensitive to injustice perpetrated against others but had no tendency to be particularly sensitive to injustice committed against themselves, a finding that HCB is associated with a morally motivated concern for others rather than a tendency to feel victimized. However, they also had elevated scores on personal vulnerability, a facet of Neuroticism, implying a heightened response to some kinds of threat, most likely social in nature. Women scored higher than men on HCB but there was only a weak and inconsistent tendency for younger people to have broad concepts, contrary to suggestions that they are a generational vanguard for creeping concepts. McGrath and Haslam (2020) showed that associations between HCB and other individual differences measures do not reflect a generalized tendency to hold more inclusive concepts.

Current research indicates that holding broad concepts of harm is associated with an array of prosocial and other-regarding dispositions, but also with a less favorable tendency to feel personally vulnerable. These findings point to HCB as having mixed blessings but give little support for the view that people with broad concepts of harm are hyper-sensitive "snowflakes." The proposed centrality of harm is supported by HCB's robust links to harm-based morality and liberalism, which is itself correlated with favoring harm and care as a primary basis for moral judgment. It is also supported by the finding that people holding a broader concept of sexism were especially likely to identify a victim of ambiguous gender discrimination and harassment as deeply harmed and morally deserving, and the perpetrators as especially blameworthy (Chan & Haslam, 2019).

Cross-sectional findings such as these accord with the theory of concept creep by showing that harm is implicated in individual differences in the inclusiveness of creeping concepts. However, they are at a substantial remove from concept creep as a historical phenomenon, which only three studies have explored directly. Wheeler et al. (2019) examined trends in the salience of the five foundations proposed by Moral Foundations Theory across the 20th century using the Google Books (Google NGram) corpus. Using word sets ("dictionaries") developed to assess the foundations, they plotted mean trajectories of the relative frequencies of words in each set from 1900 to 2007 (at the time of the study, the corpus ended in 2008). Compatible with the hypothesis that concept creep reflects a rising sensitivity to harm in the past few decades, Wheeler et al. observed a steep rise in the prominence of the Harm foundation from 1980 to 2007; a discrete rise not observed for any other foundation.

Wheeler et al.'s (2019) frequency-based findings of a rise in the cultural salience of harm do not demonstrate the semantic inflation of harm-related concepts proposed by the theory of concept creep. Two more recent contributions explore such inflation in different ways. Fabiano and

Haslam (2020) asked whether the concept of mental disorder had broadened across successive editions of the *DSM*, as critics have argued (e.g., Frances, 2013). Using meta-analytic methods to determine whether disorders are diagnosed at higher ratings when using criteria from later *DSM* editions relative to earlier ones, they found no wholesale evidence of systematic diagnostic inflation from *DSM–III* (1980) to *DSM–5* (2013), although specific conditions such as attention-deficit-hyperactivity disorder (ADHD) clearly inflated. In contrast, Vylomova et al. (2019) used two computational linguistic methods—latent semantic analysis (Sagi et al., 2009) and a prediction-based language modeling approach (Hamilton et al., 2016)—to evaluate changes in semantic breadth in a sample of five harm-related concepts (addiction, bullying, harassment, prejudice, and trauma) from 1980 to 2017. They conducted their analysis using a new corpus composed of the scientific abstracts of over 800,000 articles published in psychology journals, enabling a substantial analysis of conceptual change within academic psychology, the main crucible of semantic inflation in harm-related concepts according to the theory of concept creep. Vylomova et al. demonstrated patterns of semantic broadening in these concepts, further showing that the greatest semantic changes took place from the 1980s to the 1990s and elaborating the detail of some semantic expansions. For example, "addiction" became increasingly associated with "gaming" and "Internet" over the study period, bullying became increasingly associated with "workplace," and "trauma" became increasingly psychologized. This work indicated that the rapidly evolving computational tools of researchers in lexical semantic change can be used in large scale studies of concept creep, understood as a process of cultural dynamics.

Illustrative Analysis

To illustrate how such studies might proceed, we report a series of analyses that unpack the process of concept creep in ways that advance beyond prior work on the topic. These analyses are in some respects preliminary rather than definitive, exemplifying how more detailed future investigations, reported in greater detail, might proceed. Despite their preliminary nature, these illustrative analyses innovate in (a) examining the rising salience of harm within psychology as well as in the wider culture; (b) examining and comparing multiple text corpora; (c) using time series analyses to explore potential causal processes implicated in concept creep; and (d) presenting a new framework for examining dynamic processes involved in concept creep, and in particular how psychology articulates with the wider culture in contributing to the semantic inflation of harm. If, as concept creep theory proposes, psychology's expanding concepts of harm play a role in expanding concepts circulating outside its journals, textbooks, classrooms, and conferences, then such impacts should be detectable with appropriate data and methods. A framework for making sense of these processes in presented in Figure 1.

Figure 1
Framework for Examining Cultural Dynamics of Conceptual Change (A and B Represent Diffusion Processes; C and D Represent Language Processes)

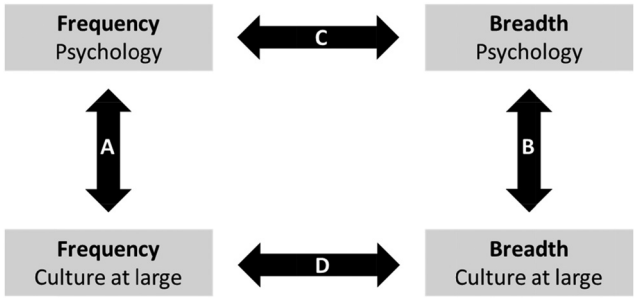

Figure 1 presents a model of possible influence processes involving the conceptual domains of psychology and the culture at large. All arrows are two-headed, indicating that influences might operate in both directions, including simultaneously. Arrow A represents processes whereby changes in the relative frequency of a concept in psychology (as an index of its prominence or salience) generate corresponding changes in the wider culture, or the reverse process whereby changes in the concept's salience in the wider culture bring about corresponding changes in its salience in psychology. The former processes involve diffusion from academic to public discourse whereas the latter involve academic discourse adopting the preoccupations of the wider society (e.g., by increased topicality or social relevance). Arrow B represents the comparable change processes involving shifts in concept meaning rather than salience. Altered concept meanings in psychology either diffuse outward into the culture at large or psychology adopts changes in concept meanings that originated elsewhere.

Whereas arrows A and B involve between-domain diffusion dynamics, arrows C and D involve within-domain language processes involving relationships between frequency and semantic change, relationships that have been addressed within computational linguistics by Hamilton et al. (2016). Arrow C represents conceptual changes occurring within psychology: changes in the salience of a concept in the field generate changes in its meaning (e.g., greater use of the concept producing semantic expansion as novel meanings arise), or changed meanings alter the frequency with which the concept is used (e.g., broadened meanings of a concept make it relevant to new contexts and thereby increase frequency of usage). Finally, arrow D represents comparable processes in the culture at large. In part of analysis that follows, an approach to examining these eight possible

influences for a selected set of harm-related concepts is demonstrated.

The analyses presented below use four massive text corpora to address three fundamental questions in the study of concept creep. First, is the cultural salience of harm rising in the wider culture in recent years, and is it also rising within the discourse of psychology? Second, are particular harm-related concepts increasing in their salience and semantic breadth within the wider culture and within psychology, as concept creep theory proposes? Third, can the framework presented in Figure 1 clarify the dynamic processes that contribute to concept creep within academic psychology and the culture at large?

Text Corpora

Four text corpora were used to examine the research questions. First, we used the recently released "English (2019)" corpus from Google Books, accessed via NGram Viewer, which updates the coverage of books from the previous Google Books endpoint of 2008. Although some concerns have been raised about the reliability of the last decade of the corpus, we used it for consistency with the prior analysis by Wheeler et al. (2019) that used the earlier Google Books version. Second, we used a new corpus of 871,340 article abstracts, originally described by Vylomova et al. (2019), which was drawn from 875 psychology journals in the E-Research and the PubMed databases for the period 1930–2019. Because abstracts are relatively sparse earlier in the corpus, and incomplete for 2018 and 2019; we restricted the analyses to abstracts of articles published between 1975 and 2017. Third, we used the Corpus of Historical American English (CoHA; Davies, 2012), which runs from the 1810s to the early 2000s and contains 400 million words from 115,000 texts evenly sampled for each decade from fiction, magazines, newspapers, and nonfiction books. Fourth and finally, we used the Corpus of Contemporary American English (CoCA; Davies, 2010), which covers the period from 1990 until 2019 and contains about 560 million words from approximately 500,000 texts evenly sampled from spoken language, TV shows, academic journals, fiction, magazines, newspapers, and blogs. For the purposes of the analyses, we combined CoHA and CoCA and extracted text for the period 1975–2017 for equivalence to the psychology abstracts corpus. From CoCA we excluded text from blogs because of the lack of temporal marks and also removed text extracted from academic journals so as not to contaminate the comparison with the psychology corpus. The psychology abstracts and CoHA/CoCA corpora were tokenized in the same fashion, with numbers and punctuation, and stop-words removed (except where noted below), and case-folding (reducing all letters to lower case) and lemmatization were carried out. Although the CoHA/CoCA corpus was a composite, it displayed no evident discontinuities in our measures of relative frequency or semantic breadth at the years where two corpora overlapped, indicating that combining them was not problematic.

Research Question 1: The Rising Salience of Harm

We first asked whether the rising salience of harm, documented by Wheeler et al. (2019) in the Google Books corpus from about 1980 to 2007, has continued since then, and whether the same rise has occurred in psychology. To address this question, we used the Google Books NGram English (2019) and psychology abstracts corpora. For comparability with the NGram corpus, the raw data (before tokenizing and removal of items as noted above) were used for the analysis of the abstracts corpus. Of the 50 terms in the Moral Foundations harm dictionary (that included some positively valenced care-related words) we used the 43 terms that were not also included in other Moral Foundations dictionaries, and in the case of terms that were word stems we selected the shortest word that could be formed from that stem. We then calculated the combined relative frequency of the 43 harm words in each year from 1975–2017 in the two corpora.

Figure 2 documents the same rise in the prominence of harm-related words in the Google Books corpus that Wheeler et al. (2019) demonstrated and shows that it continued after 2007. Overall, the summed relative frequency of the harm dictionary words increases 72% from 1975 to its peak in 2015. However, the 238% increase evident in the psychology abstracts corpus from 1975 to 2017 is markedly steeper. Despite being equally prevalent in the 1970s, the harm words proliferated in psychology article abstracts much more rapidly than they did in Google Books. This finding indicates that the rising prominence of harm replicates beyond a corpus representing the culture at large. Its precipitous increase in psychology might play a role in the broadening of harm-related concepts in the discipline. That increase in the prominence of harm may reflect a shift in the topics that academic psychology addresses, a shift in how existing topics are framed, or stylistic shifts in psychology writing—alternative explanations that have yet to be teased apart—but it represents a real increase all the same.

Research Question 2: Extending the Scope of Concept Creep

Vylomova et al. (2019) provided evidence that a sample of harm-related concepts had undergone semantic broadening in psychology article abstracts between 1980 and 2017. However, they did not provide a means of establishing whether this broadening was statistically significant, or simultaneously assess changes in the salience of the concepts. No one has yet demonstrated the semantic broadening of harm-related concepts outside of the psychology

Figure 2
Summed Relative Frequency of Harm Words in the Google Books NGram and Psychology Abstracts Corpora (1975–2017)

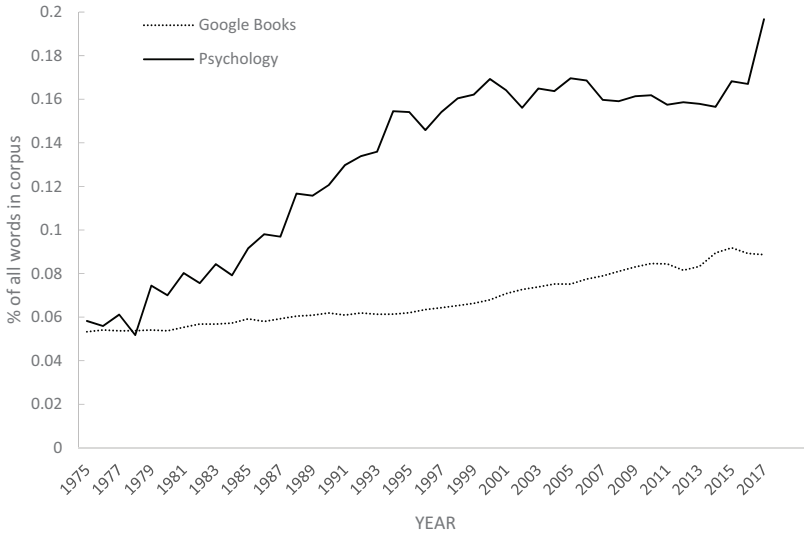

abstracts corpus. To overcome these limitations, we computed the relative frequencies and semantic breadths of a set of harm-related concepts (addiction, bullying, empathy, racism, and trauma) in the psychology abstracts corpus and in the combined CoHA/CoCA corpus. This far from exhaustive sample deliberately selected a diverse assortment of concepts with primary relevance to different fields of psychology: bullying to developmental and organizational psychology, addiction and trauma to clinical psychology, and racism to social psychology. Empathy was sampled because although concept creep research has focused on undesirable forms of harm, concept creep theory applies to all harm-related concepts, including those involving care (i.e., concern with, responsiveness to, and protection from harm). The other four concepts included examples of putatively creeping concepts relevant to clinical, developmental, and social psychology discussed by Haslam (2016) and were not intended to represent creeping concepts comprehensively.

The focus of the present analyses was on concepts rather than single words, so we selected multiple words to represent each concept (i.e., words that shared the same word stem). The words selected to represent a particular concept were each required to occur at least 50 times both in the psychology and in the combined CoHA/CoCA corpus. By this selection criterion, the concept of addiction was represented by "addict," "addiction," and "addicted"; bullying by "bully" and "bullying"; empathy by "empathy," "empathic," and "empathize"; racism but "racism" and "racist"; and trauma by "trauma," "traumatic," and "traumatize." To derive the salience data for each corpus, we computed the summed relative frequency of the words representing each of the five concepts each year as a proportion of all words in that year. Figure 3a–e presents their trajectories and their relative elevations in the two corpora. With the exception of racism, all concepts appeared at substantially higher frequencies in the psychology abstracts, consistent with expectations that harm-related concepts would be prominent in the discipline. Racism may be an exception because it is a major focus of societal concern that draws extensive attention from newspapers, TV shows, and magazines, which are a large component of the CoHA/CoCA corpus. Table 1 further shows that every concept tended to rise in salience over time in both corpora, as indicated by the positive correlations between relative frequency and year, in line with a generalized increase in the salience of harm.

To evaluate semantic breadth, we used a methodology drawn from computational linguistics (see online supplementary materials for a detailed description). In essence, the methodology has three basic steps. First, a neural network model, which represents words as numerical vectors in a multidimensional space, is trained on the text corpus (CoHA/CoCA or the psychology abstracts). These vectors, called "word embeddings," are based on the words that commonly occur in the target word's immediate context (i.e., those falling shortly before or after it in the corpus). Words with similar meanings have vectors that project at similar angles in the space. The smaller the angular displacement of two vectors, the larger the cosine of the angle, so semantic similarity can be quantified as "cosine similarity." Second, specific usages of the concept of interest (e.g., instances where the words addict, addiction, or addicted appear, for the concept of addiction) are sampled from each

Figure 3

Relative Frequencies of the Five Harm-Related Concepts in the Psychology Abstracts and General (CoHA/CoCA) Corpora

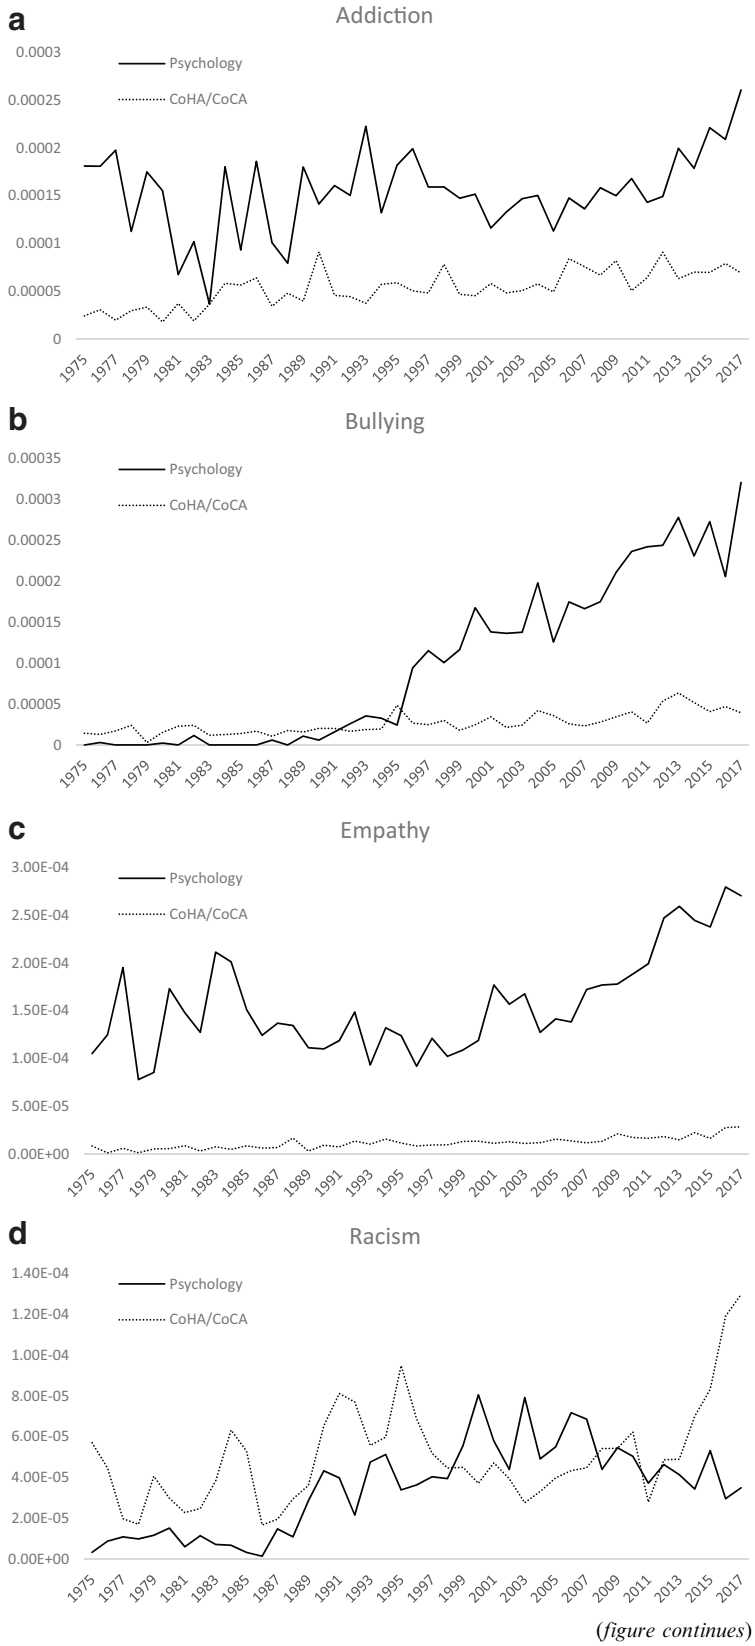

(*figure continues*)

Figure 3. (*continued*)

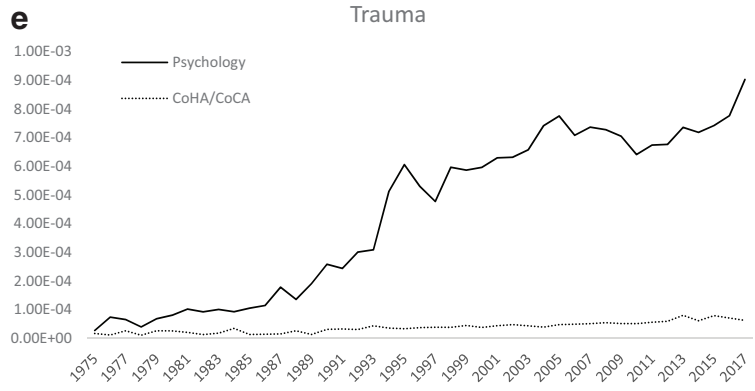

Note. CoHA = Corpus of Historical American English; CoCA = Corpus of Contemporary American English.

year of the corpus. Vectors representing each specific usage ("sentence-specific representations") are derived. Third, the cosine similarity of pairs of these specific representations in a particular year is calculated (e.g., the similarity of the vectors representing two specific usages of bullying in the year 1993), and the mean cosine similarity, across all of these pairs, is computed for each concept in each year. This mean represents the average semantic similarity of the usages of the concept in that year. Concepts with greater semantic breadth have lesser average cosine similarity because the usages of the concept appear in more diverse semantic contexts. Because semantic breadth is inversely related to mean cosine similarity, we reverse-score the latter for our index of breadth. If concepts are undergoing concept creep, this reverse-scored index should rise over time.

Table 1 reveals strong evidence for diachronic increases in the semantic breadth of the concepts. With the exception of racism (p = .063), all of the concepts broadened significantly in the psychology corpus and three of the five broadened in CoHA/CoCA. These findings provide further support for the existence of concept creep, showing that the harm-related concepts have consistently broadened over recent decades. Notably, this semantic dilation can be observed in a general corpus representing the wider culture as well as in an academic psychology corpus, and it characterizes at least one positive harm-related concept (empathy) in addition to the negatively valenced concepts that have been the focus of most previous work on concept creep.

Research Question 3: Exploring the Dynamics of Concept Creep

The preceding analyses present strong evidence that harm-related concepts have consistently risen in salience both within the general culture and, most strikingly, in academic psychology, and that with some exceptions they have also broadened their meanings in both domains. These findings justify a more sophisticated investigation of dynamic processes that may be involved in changes in concept salience and meaning. We are now in possession of all forms of time series data represented in Figure 1—semantic breadth and relative frequency in general and psychology corpora—and can examine the eight possible directional effects identified in the figure for the five harm-related concepts.

To carry out the 40 required time-series analyses, we ran cross-lagged models for each DV/IV pairing one at a time —in econometrics these are referred to as autoregressive distributed lag (ARDL) models. Each model includes as predictors time (year), the lagged DV (autoregressive effect), and the lagged IV. Using an automated procedure developed by Gluzmann and Panigo (2015), the analysis selects the best model based on the Akaike information criterion (AIC) with lags of up to 6 years for the DV and up to 7 years for the IV. Coefficients from the best-fitting model for each DV/IV pair were then examined for evidence supporting any of the influence processes schematized in Figure 1.

Table 1
Correlations Between Year and Indices of Concept Salience and Breadth

	Psychology abstracts		CoHA/CoCA	
Concept	Relative frequency	Semantic breadth	Relative frequency	Semantic breadth
Addiction	.33*	.74**	.73**	.32*
Bullying	.94**	.53**	.77**	.04
Empathy	.60**	.78**	.83**	−.18
Racism	.71**	.35	.46**	.41*
Trauma	.95**	.46**	.92**	.47**

Note. Period ranges from 26–43 years for different analyses. CoHA = Corpus of Historical American English; CoCA = Corpus of Contemporary American English.
* $p < .05$. ** $p < .01$.

For example, to test the influence process represented by the downward arrow A in Figure 1 for the concept of bullying—that is, do changes in the relative frequency (salience) of bullying in the psychology abstracts influence changes in the relative frequency of bullying in the wider culture?—we examine the two time series presented in Figure 3b. Taking the CoHA/CoCA series as the DV, we ran a regression model in which its values were predicted by (a) the previous year's value of the DV (and potentially longer time lags as well), (b) the previous year's value of the psychology abstracts IV (and potentially also longer lags), and (c) year. The critical test of the influence process is whether the relative frequency of bullying in CoHA/CoCA is predicted by (b) independently of autoregressive (a) and time (c) effects.

Twenty-nine of the 40 selected models were statistically significant overall in terms of R^2 (all analyses are reported in the online supplementary materials). Models in which the DVs involved relative frequency were more likely to be significant than those in which the DVs involved semantic breadth (18 vs. 11). These models were also more likely to yield significant time effects (10 [all positive, representing increases] vs. 5 [three positive, two negative]) and autoregressive effects (12 [nine positive, three negative] vs. 2 [both negative]), suggesting that the breadth time series were less predictable and structured. The abundance of autoregressive effects provides evidence of an endogenous, self-reinforcing dynamic within psychology (seven positive frequency effects) and the wider culture (one positive effect) for increasing attention to the harm-related concepts.

Significant lagged effects of the IVs emerged in 11 of the 40 analyses. These effects, organized in terms of the influence processes presented in Figure 1, are summarized in Table 2. The table indicates that at least one concept demonstrated each of the four proposed processes (A to D), that six of the eight possible influence directions were supported by at least one concept, and that every concept was involved in at least one significant effect. The analyses support the plausibility of the influence processes that we proposed, even if significant effects were only moderately common, and they motivate more detailed investigations of specific conceptual changes. We repeated all of the 40 analyses using detrended (i.e., first-differenced) versions of the IVs and DVs, and these analyses yielded the identical number of significant predictor effects (i.e., 11), six of which replicated the original findings, with several others flipping from marginal to significant or the reverse. Thus, the evidence for the plausibility of the influence processes appears to be robust when the variables are detrended.

Table 2 suggests that evidence for the proposed language processes (C and D) was stronger than evidence for the diffusion processes (A and B), with eight and three significant effects, respectively. The rising salience (relative frequency) of a concept had lagged effects on the breadth of its meaning, and changes in concept breadth had lagged effects on changes in concept salience. Effects in the two directions were equally common, and these effects were found in both the psychology and general corpora. Diffusion effects were scarce. There was some evidence that shifts in the frequency of mentions of racism in psychology predicted subsequent shifts in the opposite direction in its salience in the wider culture, as represented by CoHA/CoCA, and some evidence that changes in the semantic breadth of empathy and racism within the wider culture generated corresponding changes in the breadth of these concepts within psychology. Nevertheless, evidence for diffusion processes is patchy.

In summary, the time series analyses indicate that diachronic changes in the salience and semantic breadth of harm-related concepts can be modeled by a combination of time, endogenous, and dynamic influence processes. These processes link different dimensions of concepts (salience and semantic breadth) and different contexts in which they are used (academic psychology and everyday text). However, evidence that conceptual expansions initiated in the academy diffuse into the ambient culture, consistent with the view that concept creep tends to originate in psychology, is weak. There are several possible reasons for this lack of evidence. Actual diffusion effects may operate at a time scale that the data are not suitable to detect. For example, effects of changed concept salience in psychology on

Table 2
Summary of Significant Lagged Predictor Effects From the Time Series Analyses

Process	Domain	Direction (IV → DV)	Significant effects
Diffusion processes	A. Frequency	Psychology → culture	Racism (-)
		Culture → psychology	—
		Psychology → culture	—
	B. Breadth	Culture → psychology	Empathy, racism
Language processes	C. Psychology	Frequency → breadth	Empathy (-), racism, trauma
	Breadth → frequency	Addiction, empathy (-)	
	D. Culture	Frequency → breadth	Bullying (-)
		Breadth → frequency	Bullying, trauma (-)

Note. IV = independent variable; DV = dependent variable.

changed salience in the wider culture might be relatively immediate (e.g., within-year effects due to rapid science communication). However, the theory of concept creep may also overstate the proximal role of psychology in driving the inflation of harm-related concepts in society. Any such role may be weak or limited to particular concepts. Distal explanations that view concept creep as the consequence of a rising sensitivity to harm or as an adaptive response to reductions in material hardship may remain tenable, but proximal explanations that attribute a major causal role to psychology may be questionable.

Several caveats are for accepting this verdict, however. First, as some analyses interrogated time series as short as 26 years—because particular concepts were not represented in sufficient numbers early in one or both corpora—and none examined series longer than 43 years, they were arguably underpowered. To some degree this limitation is intrinsic to the study of diachronic processes that unfold over a few decades and that cannot readily be broken into smaller temporal units for operational or measurement reasons. However, the shortness of the time series not only reduces the capacity to detect modest effects but also rules out the investigation of more complex quantitative models. Second, the chosen method for evaluating semantic breadth as cosine (dis)similarity of word embeddings is just one of many that have been developed in recent years (e.g., Dubossarsky et al., 2017; Tahmasebi et al., 2018), and may not capture some forms of concept creep. By assessing semantic breadth as the diversity of semantic contexts in which a concept appears, it may be well suited to appraising horizontal creep. However, it may be incapable of detecting vertical creep, where concepts come to refer to less severe phenomena in an existing context rather than attaching to qualitatively new contexts. Further studies using different corpora, more fine-grained or longer duration time series, and alternative methods for evaluating concept breadth are needed to address these and related concerns.

Conclusions

Investigation of our three research questions helps to clarify the empirical status of concept creep. Analysis of the first question reveals that the rising salience of harm-related words, previous observed in a general corpus up to 2007 (Wheeler et al., 2019) has continued its upward trajectory to the present day. The significance of this rise, and the steadily increasing cultural prominence of harm that it implies, have yet to be determined. Our new finding that the same rise appears in our corpus of psychology abstracts, where it is three times steeper, is noteworthy in that context. Harm permeates recently published psychology abstracts substantially more than it does a general corpus of English language books, despite being equally represented in the two corpora in the 1970s. This steeper rise on psychology is mirrored in the typically higher salience and steeper increase of the five putatively creeping concepts in psychology relative to another general corpus in Figure 3 and the typically stronger increases in the semantic breadth of these concepts in psychology suggested by the correlations reported in Table 1.

Analysis of our second question also extends previous research on concept creep in new directions. Our five harm-related concepts all rose in relative frequency over the study period both in psychology abstracts and in a carefully curated general corpus (CoHA/CoCA), and they were generally much more frequent in the former. The demonstration that empathy shows the same rising trajectory as other harm-related terms is consistent with the claim that what is rising is the salience of harm in a broad sense that encompasses positively valued ideas of care, safety, and protection rather than in a narrowly negative sense. Concept creep, the progressive broadening of harm-related concepts, was also demonstrated consistently, shown in almost all cases in the psychology corpus and also, for the first time but somewhat less robustly, in a general corpus. These findings confirm that concept creep, and the rising prominence of harm-related concepts that is linked to it, is occurring both within psychology and in the culture at large.

Our final set of analyses provide moderate support for our new framework of dynamic processes involved in concept creep. They found many examples of temporal trends consistent with rises in the salience and breadth of harm-related concepts, and of endogenous tendencies for selected concepts to increase in salience or, less commonly, in breadth in self-reinforcing ways, both within psychology and without. Many examples of lagged effects consistent with causal influences operating between semantic breadth and relative frequency, or more rarely between corpora, were also obtained. These illustrative analyses offered positive plausibility tests of all of our proposed dynamic influence processes. However, the slender evidence for lagged effects of changes in the psychology corpus on changes in the general corpus undermines any strong claim about psychology playing a dominant role in disseminating inflated concepts of harm.

Our analyses raise as many questions as they answer. It is unclear why psychology has become so much more harmified (Schein & Gray, 2016) in recent decades. Harm-based morality is associated with a liberal political orientation (Graham et al., 2009) and the rising salience and broadening definition of harm within academic psychology might be tied to the increasingly liberal values in the field (Duarte et al., 2015; Redding, 2001) or instead to shifts in topics that are not intrinsically politicized. Whether concept creep is at all specific to harm-related concepts, and even whether the most appropriate conceptualization of harm is narrow, as in Moral Foundations Theory, or broad enough to be the grounding of all morality, as in the Theory of Dyadic Morality (Schein & Gray, 2018), remains to be seen. The

cultural boundaries of concept creep are also unknown. Weaker evidence for it in the general corpus may indicate that broadened concepts of harm may be held only by particular segments of the public. Whether creep can be demonstrated in cultural settings beyond the Anglosphere is another urgent question requiring further study.

Although the evidence for psychology's role in promoting concept creep is weak, there is ample reason to believe that concept creep and an intensifying focus on harm are occurring within our field. Understanding what is driving the rising salience and expanding meanings of harm in the discipline, and the cultural, societal, and scientific implications of these trends, should be a research priority. Answering these questions will require a psychology of cultural dynamics that can embrace and advance the study of meaning change. Such a psychology offers an ideal vantage point for making sense of why we are increasingly in harm's way.

References

Alesina, A., Giuliano, P., & Nunn, N. (2013). On the origins of gender roles: Women and the plough. *The Quarterly Journal of Economics*, *128*(2), 469–530. https://doi.org/10.1093/qje/qjt005

Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. University of Chicago Press.

Campbell, B., & Manning, J. (2018). *The rise of victimhood: Microaggressions, safe spaces, and the new culture wars*. Palgrave Macmillan. https://doi.org/10.1007/978-3-319-70329-9

Cavalli-Sforza, L. L., & Feldman, M. W. (1981). *Cultural transmission and evolution*. Princeton University Press.

Chan, J., & Haslam, N. (2019). Broad concepts of sexism predict polarized moral judgments of victims and perpetrators. *Personality and Individual Differences*, *150*, 109488. https://doi.org/10.1016/j.paid.2019.06.031

Chiu, C.-Y., & Hong, Y.-Y. (2013). *Social psychology of culture*. Psychology Press. https://doi.org/10.4324/9781315782997

Davies, M. (2010). The corpus of contemporary American English as the first reliable monitor corpus of English. *Literary and Linguistic Computing*, *25*(4), 447–464. https://doi.org/10.1093/llc/fqq018

Davies, M. (2012). Expanding horizons in historical linguistics with the 400-million word Corpus of Historical American English. *Corpora*, *7*(2), 121–157. https://doi.org/10.3366/cor.2012.0024

De Vos, J. (2013). *Psychologization and the subject of late modernity*. Palgrave Macmillan. https://doi.org/10.1057/9781137269225

Duarte, J. L., Crawford, J. T., Stern, C., Haidt, J., Jussim, L., & Tetlock, P. E. (2015). Political diversity will improve social psychological science. *Behavioral and Brain Sciences*, *38*, e130. https://doi.org/10.1017/S0140525X14000430

Dubossarsky, H., Weinshall, D., & Grossman, E. (2017). Outta control: Laws of semantic change and inherent biases in word representation models. *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing* (pp. 1136–1145). Association for Computational Linguistics.

Fabiano, F., & Haslam, N. (2020). Diagnostic inflation in the *DSM*: A meta-analysis of changes in the stringency of psychiatric diagnosis from *DSM–III to DSM–5*. *Clinical Psychology Review*, *80*, 101889. https://doi.org/10.1016/j.cpr.2020.101889

Frances, A. (2013). *Saving normal: An insider's revolt against out-of-control psychiatric diagnosis, DSM–5, big pharma, and the medicalization of ordinary life*. William Morrow.

Furedi, F. (2004). *Therapy culture: Cultivating vulnerability in an uncertain age*. Routledge.

Gluzmann, P., & Panigo, D. (2015). Global search regression: A new automatic model-selection technique for cross-section, time-series, and panel-data regressions. *The Stata Journal*, *15*(2), 325–349. https://doi.org/10.1177/1536867X1501500201

Graham, J., Haidt, J., & Nosek, B. A. (2009). Liberals and conservatives rely on different sets of moral foundations. *Journal of Personality and Social Psychology*, *96*(5), 1029–1046. https://doi.org/10.1037/a0015141

Graham, J., Nosek, B. A., Haidt, J., Iyer, R., Koleva, S., & Ditto, P. H. (2011). Mapping the moral domain. *Journal of Personality and Social Psychology*, *101*(2), 366–385. https://doi.org/10.1037/a0021847

Greenfield, P. M. (2013). The changing psychology of culture from 1800 through 2000. *Psychological Science*, *24*(9), 1722–1731. https://doi.org/10.1177/0956797613479387

Grossmann, I., & Varnum, M. E. W. (2015). Social structure, infectious diseases, disasters, secularism, and cultural change in America. *Psychological Science*, *26*(3), 311–324. https://doi.org/10.1177/0956797614563765

Hacking, I. (1995). The looping effect of human kinds. In D. Sperber, D. Premack, & A. J. Premack (Eds.), *Causal cognition: A multidisciplinary debate* (pp. 351–383). Oxford University Press.

Hamamura, T. (2012). Are cultures becoming individualistic? A cross-temporal comparison of individualism-collectivism in the United States and Japan. *Personality and Social Psychology Review*, *16*(1), 3–24. https://doi.org/10.1177/1088868311411587

Hamamura, T., & Xu, Y. (2015). Changes in Chinese culture as examined through changes in personal pronoun usage. *Journal of Cross-Cultural Psychology*, *46*(7), 930–941. https://doi.org/10.1177/0022022115592968

Hamilton, W. L., Leskovec, J., & Jurafsky, D. (2016). Diachronic word embeddings reveal statistical laws of semantic change. *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 1489–1501). https://doi.org/10.18653/v1/P16-1141

Haslam, N. (2016). Concept creep: Psychology's expanding concepts of harm and pathology. *Psychological Inquiry*, *27*(1), 1–17. https://doi.org/10.1080/1047840X.2016.1082418

Haslam, N., Dakin, B. C., Fabiano, F., McGrath, M. J., Rhee, J., Vylomova, E., Weaving, M., & Wheeler, M. A. (2020). Harm inflation: Making sense of concept creep. *European Review of Social Psychology*, *31*(1), 254–286. https://doi.org/10.1080/10463283.2020.1796080

Heine, S. J. (2020). *Cultural psychology*. Norton.

Inglehart, R. F. (2008). Changing values among Western publics from 1970 to 2006. *West European Politics*, *31*(1–2), 130–146. https://doi.org/10.1080/01402380701834747

Kashima, Y. (2001). Culture and social cognition: Towards a social psychology of cultural dynamics. In D. Matsumoto (Ed.), *Handbook of culture and psychology* (pp. 325–360). Oxford University Press

Kashima, Y. (2008). A social psychology of cultural dynamics: Examining how cultures are formed, maintained, and transformed. *Social and Personality Psychology Compass*, *2*(1), 107–120. https://doi.org/10.1111/j.1751-9004.2007.00063.x

Kashima, Y., Bain, P. G., & Perfors, A. (2019). The psychology of cultural dynamics: What is it, what do we know, and what is yet to be known? *Annual Review of Psychology*, *70*, 499–529. https://doi.org/10.1146/annurev-psych-010418-103112

Kashima, Y. (2014). How can you capture cultural dynamics? *Frontiers in Psychology*, *5*, 995. https://doi.org/10.3389/fpsyg.2014.00995

Kesebir, S., & Kesebir, P. (2017). A growing disconnection from nature is evident in cultural products. *Perspectives on Psychological Science*, *12*(2), 258–269. https://doi.org/10.1177/1745691616662473

Lamberson, P. J. (2016). Diffusion in networks. In Y. Bramoullé, A. Galeotti, & B. Rogers (Eds.), *The Oxford handbook of the economics of networks* (pp. 479–503). Oxford University Press.

Lasch, C. (1979). *The culture of narcissism: American life in an age of diminishing expectations*. Norton.

Levari, D. E., Gilbert, D. T., Wilson, T. D., Sievers, B., Amodio, D. M., & Wheatley, T. (2018). Prevalence-induced concept change in human judgment. *Science*, *360*(6396), 1465–1467. https://doi.org/10.1126/science.aap8731

Lukianoff, G., & Haidt, J. (2018). *The coddling of the American mind: How good intentions and bad ideas are setting up a generation for failure*. Penguin.

McGrath, M. J., & Haslam, N. (2020). Development and validation of the Harm Concept Breadth Scale: Assessing individual differences in harm inflation. *PLoS ONE*, *15*(8), e0237732. https://doi.org/10.1371/journal.pone.0237732

McGrath, M. J., Randall-Dzerdz, K., Wheeler, M. A., Murphy, S., & Haslam, N. (2019). Concept creepers: Individual differences in harm-related concepts and their correlates. *Personality and Individual Differences*, *147*, 79–84. https://doi.org/10.1016/j.paid.2019.04.015

Mesoudi, A. (2011). *Cultural evolution: How Darwinian theory can explain human culture and synthesize the social sciences*. University of Chicago Press. https://doi.org/10.7208/chicago/9780226520452.001.0001

Michel, J.-B., Shen, Y. K., Aiden, A. P., Veres, A., Gray, M. K., Pickett, J. P., Hoiberg, D., Clancy, D., Norvig, P., Orwant, J., Pinker, S., Nowak, M. A., & Aiden, E. L., & The Google Books Team. (2011). Quantitative analysis of culture using millions of digitized books. *Science*, *331*(6014), 176–182. https://doi.org/10.1126/science.1199644

Oishi, S., Graham, J., Kesebir, S., & Galinha, I. C. (2013). Concepts of happiness across time and cultures. *Personality and Social Psychology Bulletin*, *39*(5), 559–577. https://doi.org/10.1177/0146167213480042

Pinker, S. (2011). *The better angels of our nature: Why violence has declined*. Viking.

Pinker, S. (2018). *Enlightenment now: The case for reason, science, humanism, and progress*. Penguin.

Redding, R. E. (2001). Sociopolitical diversity in psychology. The case for pluralism. *American Psychologist*, *56*(3), 205–215. https://doi.org/10.1037/0003-066X.56.3.205

Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press.

Sagi, E., Kaufmann, S., & Clark, B. (2009). *Semantic density analysis: Comparing word meaning across time and phonetic space*. Proceedings of the Workshop on Geometrical Models of Natural Language Semantics *(pp. 104–111). Association for Computational Linguistics.*

Santos, H. C., Varnum, M. E. W., & Grossman, I. (2017). Global increases in individualism. *Psychological Science*, *28*(9), 1128–1139.

Schein, C., & Gray, K. (2016). Moralization and harmification: The dyadic loop explains how the innocuous becomes harmful and wrong. *Psychological Inquiry*, *27*(1), 62–65. https://doi.org/10.1080/1047840X.2016.1111121

Schein, C., & Gray, K. (2018). The theory of dyadic morality: Reinventing moral judgment by redefining harm. *Personality and Social Psychology Review*, *22*(1), 32–70. https://doi.org/10.1177/1088868317698288

Schulz, J. F., Bahrami-Rad, D., Beauchamp, J. P., & Henrich, J. (2019). The Church, intensive kinship, and global psychological variation. *Science*, *366*(6466), eaau5141. https://doi.org/10.1126/science.aau5141

Sperber, D. (1985). Anthropology and psychology: Towards and epidemiology of representations. *Man*, *20*(1), 73–89. https://doi.org/10.2307/2802222

Sperber, D. (1996). *Explaining culture: A naturalistic approach*. Blackwell.

Sunstein, C. R. (2018). *The power of the normal*. SSNR. https://doi.org/10.2139/ssrn.3239204

Tahmasebi, N., Borin, L., & Jatowt, A. (2018). Survey of computational approaches to lexical semantic change. *arXiv*. https://doi.org/arxiv.org/abs/1811.06278

Twenge, J. M., & Campbell, W. K. (2008). Increases in positive self-views among high school students: Birth-cohort changes in anticipated performance, self-satisfaction, self-liking, and self-competence. *Psychological Science*, *19*(11), 1082–1086. https://doi.org/10.1111/j.1467-9280.2008.02204.x

Twenge, J. M., Campbell, W. K., & Gentile, B. (2012). Increases in individualistic words and phrases in American books, 1960–2008. *PLoS ONE*, *7*(7), e40181. https://doi.org/10.1371/journal.pone.0040181

Twenge, J. M., & Foster, J. D. (2010). Birth cohort increases in narcissistic personality traits among American college students, 1982–2009. *Social Psychological & Personality Science*, *1*(1), 99–106. https://doi.org/10.1177/1948550609355719

Varnum, M. E. W., & Grossmann, I. (2017). Cultural change: The how and the why. *Perspectives on Psychological Science*, *12*(6), 956–972. https://doi.org/10.1177/1745691617699971

Vylomova, E., Murphy, S., & Haslam, N. (2019). Evaluation of semantic change of harm-related concepts in psychology. *Proceedings of the 1st International Workshop on Computational Approaches to Historical Language Change* (pp. 29–34). Association for Computational Linguistics. https://doi.org/10.18653/v1/W19-4704

Wheeler, M. A., McGrath, M. J., & Haslam, N. (2019). Twentieth century morality: The rise and fall of moral concepts from 1900 to 2007. *PLoS ONE*, *14*(2), e0212267. https://doi.org/10.1371/journal.pone.0212267

Wolff, P., Medin, D. L., & Pankratz, C. (1999). Evolution and devolution of folkbiological knowledge. *Cognition*, *73*(2), 177–204. https://doi.org/10.1016/S0010-0277(99)00051-7

Received August 16, 2020
Revision received February 2, 2021
Accepted March 2, 2021 ■