A bibliometric evaluation of the impact of theories of consciousness in academia and on social media

Andy Wai Kan Yeung¹, Cody A. Cushing², Alan L. F. Lee^{3,4}

1. *Oral and Maxillofacial Radiology, Applied Oral Sciences and Community Dental Care, Faculty of Dentistry, The University of Hong Kong*
2. *Department of Psychology, UCLA*
3. *Department of Applied Psychology, Lingnan University, Hong Kong*
4. *Wofoo Joseph Lee Consulting and Counselling Psychology Research Centre, Lingnan University, Hong Kong*

Correspondence:

Alan L. F. Lee

Email: ALANLEE@LN.EDU.HK

Phone: +852 2616 7419

DATA AVAILABILITY STATEMENT

The data that support the findings of this study are openly available in OSF at <https://osf.io/x685r/>.

ACKNOWLEDGEMENTS

We would like to thank Matthias Michel and Mouslim Cherkaoui for their constructive comments during the draft of this paper.

ABSTRACT

Consciousness science faces both opportunities and challenges as it has been maturing as a discipline of science in recent years. Some popular theories of consciousness have been described as unfalsifiable and some as “overpromoted” in the media. To objectively evaluate the state of consciousness science as a field, we analyzed the bibliometric data for five major theories of consciousness, namely, Global workspace (GWT), Higher-order (HOT), Integrated information (IIT), Local recurrent (LRT), and Quantum theories (QT). We analyzed data for both academic publications and citations (2004-2019) and Twitter activity data (2011-2019). We found that the theories varied in terms of their impact in academia and on social media. Specifically, IIT was found to have the highest and fastest-growing raw counts of publication and citation, but have lower per-publication citation counts and proportion of citations with empirical support/contradiction as evaluated by a machine-learning algorithm. The data on social media revealed that IIT was the “most-tweeted” theory, though, together with QT, Twitter mentions of their core papers mostly came from the general public instead of from professionals. Our findings suggest that a theory’s impact quality could be related to its degree of media promotion.

INTRODUCTION

The field of consciousness science has been going through structural changes in recent years, gradually turning from a “mystery” about subjective experiences to a maturing discipline of science. Such changes have led to not only opportunities for the field, but also challenges that could hinder development (Michel et al., 2019). These challenges include the difficulty for projects related to consciousness to obtain funding, difficulty for consciousness researchers to find related academic jobs, and, most importantly, the perceived lack of scientific rigor in consciousness research (Michel et al., 2018).

In particular, given the wide scope of phenomena that consciousness science as a discipline entails, one problem for the field is that theories for consciousness could generate claims that are not empirically testable. Recently, some researchers have argued that it is impossible to empirically falsify certain theories or types of theories of consciousness (Hanson & Walker, 2020, Doerig et al., 2019, but see Kleiner (2020) and Tsuchiya et al. (2019) for replies). As Michel et al. (2019) argued, if they were indeed unfalsifiable, attributing funding for work on these theories could harm the field as a whole in the long run, especially when such funding opportunities could have supported the development of other more empirically testable theories.

The problem could be exacerbated by “overpromotion” in social media (Michel et al., 2018), which has been demonstrated to have an effect on the download rate of an academic article (Botting et al., 2016). When an unfalsifiable but popular theory is heavily promoted, it could lead to growth in sheer quantity (e.g., number of publications,

citations, and social media activity) but lack of quality (e.g., empirical evidence that helps the theory or the field as a whole develop).

However, many of the above claims are either theoretical (e.g., Doerig et al., 2019; Hanson & Sara, 2020) or subjective (as a survey on researchers, Michel et al., 2018).

The present study aimed at evaluating the extent to which the above statements are correct for consciousness science with a bibliometric approach. Analyzing bibliometric data provides more objective and concrete data to describe the impact of a theory. In addition, we extracted Twitter activity as a measurement of each theory's impact in social media. By comparing the theories in terms of their impact both in academia and in social media, we can then evaluate the development of consciousness science as a field by directly addressing certain claims above, e.g., that some theories have been “overpromoted” in social media (Michel et al., 2018).

We focused on five major theories of consciousness: global workspace theory (GWT; e.g., Baars, 1988; Dehaene & Naccache, 2001; Dehaene, Kerszberg, & Changeux, 1998), higher-order theory (HOT; e.g., Rosenthal, 1986; 2005; Lau & Rosenthal, 2011; Brown, Lau, & LeDoux, 2019), integrated information theory (IIT; e.g., Tononi, 2004; Tononi, 2008; Oizumi, Albantakis, Tononi, 2014), local recurrence theory (LRT; e.g., Lamme & Roelfsema, 2000; Lamme, 2003; 2006), and quantum theory (QT; e.g., Hameroff & Penrose, 1996; 2014; Ekert, Jozsa, Penrose, & Stuart, 1998). As a preliminary survey on the Internet, we did a Google Search under the “News” tab and found substantial variability in the number of results across the theories (GWT: 120, HOT: 9, IIT: 331, LRT: 0, and QT: 1,350). We extracted bibliometric and Twitter data for

these five theories over, respectively, the past 16 years and 9 years. We then evaluated the extent to which the theories varied in their impact. Our analysis covered three main areas: yearly trend in academic impact, scope and depth of academic impact, and impact on social media.

METHODS

The electronic literature database, Scopus, was searched in late April, 2020. Five search strings, one for each of the five theories, were devised to search for all fields of the indexed publications. No restriction was placed on any other fields, such as publication type, year, etc.

For global workspace theory (GWT), the search string was:

ALL("global workspace" OR "global neuronal workspace").

For integrated information theory (IIT), the search string was:

(ALL("integrated information" AND conscious*) OR ALL((information W/2
integrat*) W/10 conscious*) OR ALL("integrated information theor*"))).

For higher order theory (HOT), the search string was:

(ALL(("higher order" PRE/3 (theor* OR model* OR approach*)) AND conscious*)
OR ALL(("higher order thought" OR "higher order perception") AND conscious*)
AND NOT ALL("higher order factor*" OR "factor analysis"))).

For local recurrent theory (LRT), the search string was:

(ALL("local recurr*" PRE/1 theor*) OR ALL("recurrent processing") AND NOT
TITLE-ABS-KEY("deep learning"))

For quantum theory (QT), the search string was:

((ALL(quantum W/10 (microtubule* OR "objective reduction" OR neuron*)) AND
ALL(consciousness))) OR (ALL("quantum mind*" OR "quantum conscious*"))

The basic bibliometric data, such as the publication counts of authorship, country, and journal were obtained by the “Analyze search results” function of Scopus. The annual citation counts were obtained by the “View citation overview” function.

Because the numbers returned by Scopus were cumulative up to a specific survey year, for each metric, the annual count for a specific year was obtained by subtracting the cumulative count of the previous year from the cumulative count of the chosen year (i.e., annual count of 2010 = cumulative count of 2010 - cumulative count of 2009).

When computing the citation count per publication for a specific year, we divided the cumulative citation count by the cumulative publication count (e.g., citation count per publication of 2010 = cumulative citation count in 2010 / cumulative publication count in 2010). This way of calculating per-publication citation allows us to take into consideration all citations across all publications published up to a certain year, which could be the case had we used annual counts instead of cumulative counts.

The basic bibliometric data of the papers citing the identified papers were obtained by clicking “View cited by” and then “Analyze search results” functions. For papers for each theory, their DOIs and PMIDs were exported and loaded into Altmetric Explorer to assess their Twitter mentions. Meanwhile, the full records were exported in csv format and loaded into CRExplorer to identify the references most cited by the analyzed papers for each theory. Since some reference information might be missing or inaccurate, the default data cleaning processing (disambiguation) in CRExplorer was conducted.

Briefly, the Levenshtein similarity values of the first authors’ last names and the journal titles were computed. If the weighted average of the two similarity values (in a ratio of

2:1) was greater than 0.75, then the concerned references were considered as the same and their citation data merged. Manual inspection was performed to separate non-duplicate references.

For each theory, from the resultant reference list in descending order of times cited by the analyzed papers, the three most cited relevant references formed the core paper collection, and were further analyzed (see supplementary materials for details of these core papers). For these core papers, the citations made to them were analyzed with scite.ai (<https://scite.ai>), a website providing a free service that classifies the citations into supporting, mentioning, and contradicting contexts with deep-learning algorithm machine-learning model. Their way of classifying citations, called “Smart citations”, is currently adopted by Europe PubMed Central (Europe PMC) and displayed next to traditional citation data.

We analyzed three aspects of impact for the five theories. First, we extracted the publication and citation data over the past 16 years (2004-2019) to analyze the trend in academic impact for each theory. Specifically, for every survey year, we extracted the number of all publications for each theory, the number of all citations for each theory, and the number of citations for the core papers of each theory. We also computed the number of citations per publication for each theory over the survey period.

Second, we focused on the latest data obtained at the time of the survey to evaluate the scope and depth of academic impact. We extracted the publishing disciplines (i.e., the disciplines of the journals in which research articles were published for each theory) and the citing disciplines (i.e., the disciplines of the journals in which the citing articles were

published for each theory). We also compared the journal impact factors (Clarivate, 2020) of the publishing journals, the number of citations per publication, and the proportion of cited publications (i.e., publications with at least one citation) across the theories. Furthermore, we attempted to evaluate the quality of citations by comparing the contexts in which the theories were cited based on the classification results from scite.ai (Khamisi, 2020).

Third, we evaluated the theories' impact on social media. We extracted the trend of Twitter activity over a 9-year period (2011-2019), in which we obtained the number of Twitter mentions for all publications under each theory. Similar to the citation data, we analyzed the number of Twitter mentions per publication and the proportion of "tweeted" publications (i.e., the publications with at least one Twitter mention). Finally, we focused on the core papers for each theory and categorized the professions of the Twitter users who mentioned the core papers on Twitter.

RESULTS

1. Trend of academic impact

We estimated the annual growth rate (r) by fitting the following function to the data, which assumes a constant annual growth rate across years

$$C_t = C_0(1 + r)^t$$

where C_t is the predicted count for a specific year (e.g., 2017), C_0 is the count of the first year in the period (i.e., set at 2004 for the period 2004-2019, except for the analysis for the citation counts of the core papers, which was set at 2005, because the core papers for IIT had zero citations in 2004 and could not be set as C_0), r is the compound annual growth rate, and t is the number of years elapsed. The model was fitted with all data within the time period by theory.

For each metric, we performed Bonferroni correction on 10 pairwise, post-hoc comparisons across the five theories.

Figure 1. Bibliometric data obtained from Scopus over years. Top row: number of publications returned by Scopus based on the theory-specific query terms. Publications related to IIT (thin solid line on the left panel) have grown at the fastest pace among all theories (middle bar on the right panel). Second row from top: number of citations returned by Scopus. For each theory, it is the total number of citations validated by Scopus that all the theory's publications have accrued in each survey year. GWT and LRT top the list, but the IIT, having the fastest-growing citation counts, has caught up in recent years. Third row from top: total number of citations for the three core papers; obtained by similar means as for the number of citations. The core papers of IIT have the highest growth rate in citation counts. Bottom row: citation counts per publication. For each year, the number was computed by dividing the cumulative citation counts up to that year by the cumulative publication counts up to the same year. Despite the high growth rates in the raw publication and citation counts, the per-publication citation counts for IIT has the slowest growth rate among all theories. Left column: yearly trends of the respective metrics. Right column: Estimated annual growth rates for respective metrics (Error bars represent 95% C.I.). Only significant (Bonferroni-corrected) pairwise differences are annotated above the bars.

1.1 Publication counts

Publication counts (Figure 1, top-left panel) for IIT have increased rapidly from 2004 to 2020. It has risen from bottom among the five theories in 2004 to the top together with GWT in 2019. The relative ranks among the other four theories have remained fairly constant in the past 10 years, with GWT at the top, followed by LRT and then by QT. QT used to be at the top of the list, but was taken over by GWT and LRT around 2008. HOT has been at the bottom on publication counts since IIT took over it in around 2010. Publication growth was the fastest for IIT (28.52%), followed by HOT (17.40%), then by GWT (13.91%), then by LRT (9.40%), and was slowest for QT (5.88%). The growth rates significantly differed from one another across all theories ($t(30) > 3.30$, Bonferroni-corrected $p < .02$).

1.2 Citation counts

For raw citation counts (Figure 1, second from top, left), GWT has been at the top throughout the survey years from 2004 to 2019. LRT has been at the second place for almost the entire period, until it was caught up by and had similar citation counts as IIT since 2018. IIT emerged as the third most-cited theory in late 2000's, and has been growing at the fastest pace in citation counts. Citation counts for HOT have been growing steadily, and over took QT to become the fourth most-cited theory in 2013. For the entire survey period (2004-2019), growth in citation counts was the fastest for IIT (31.05%), closely followed by HOT (27.72%). The growth in citations for LRT (21.96%) was the third among all five, followed by GWT (17.80%). QT was found to have the

slowest growth (15.30%). Annual growth rate in citation counts was significantly different across all theories ($t(30) > 3.30$, Bonferroni-corrected $p < .02$).

1.3 Citation counts for the core papers for each theory

When compared with the raw citation counts over all publications, the yearly trends of citation counts for the core (i.e., three most-cited) papers were not exactly the same.

The core papers of GWT had been having the most number of citations across the theories until it was taken over by IIT in 2019. Citation counts for GWT core papers peaked at 2012 (around 300 citations in that year) and appear to be decreasing since then. The citation counts for LRT's core papers show a similar trend as those for GWT's, except that it peaked at around 200 in 2009, and was taken over by the citation counts of IIT's core papers in 2016. It has stayed at the third place since then. The core papers of IIT, started with zero citations at the bottom in 2004, have been receiving an increasing number of citations over the years, until it became the most cited ones in 2019. HOT's core papers citations have grown steadily, though still lower than that for IIT's and QT's. The core papers of QT had not grown much in citation counts during 2005-2013, but have been catching since 2013 to be the second-fastest growing in citation counts.

The estimated annual growth rate of IIT's core papers (31.87%) was the highest among all theories, followed by QT (22.64%). Growth rate of HOT's core papers was at the third place (14.16%). The core papers of LRT (6.49%) and GWT (4.00%) were the slowest growing ones among all theories. Except between LRT and GWT ($t(28) = 1.66$, Bonferroni-corrected $p = .54$) and between IIT and QT ($t(28) = 2.53$, Bonferroni-corrected $p = .085$), growth rate in citations for core papers differed

significantly across all other pairs of theories ($t(28) > 3.09$, Bonferroni-corrected $p < .03$).

1.4 Citation counts per publication

GWT publications have been most cited on average for almost the entire survey period until LRT caught up in 2019 (both around 40 citations per publication). LRT, starting at its third place in 2004, took over IIT in 2008 and has been growing rapidly in citation counts per publication. IIT has remained at the third place since 2008, with relatively slow growth despite its high growth rate in raw publication and citation counts. QT had the fourth most citation counts per publication, but was surpassed by HOT in 2011, and had been at the bottom of the list since then. HOT has remained at the fourth place since around 2012.

Over the entire survey period, LRT has the fastest growth (13.38%), followed by HOT (10.75%), then by QT (7.76%) and then by GWT (6.94%). IIT has the slowest growth among all five theories (3.92%). These growth rates differed significantly for all pairs of comparisons among the five theories ($t(30) > 3.45$, Bonferroni-corrected p values $< .01$).

2. Scope and depth of academic impact

2.1 Publication and citation counts by disciplines

Figure 2. Left: Proportions of publications in journals that belong to the behavioral-sciences disciplines across theories. Right: Proportions of citations coming from journals that belong to the behavioral-sciences disciplines across theories. Error bars represent the 95% confidence intervals for the mean proportions for behavioral sciences (i.e., neuroscience + psychology). Only significant pairwise differences are marked above the bars (Bonferroni-corrected). Note the similarity between publications and citations in the patterns of results across the theories. See Figure S1 in

Supplementary Materials for a full breakdown by disciplines for publication and citation outlets.

For each theory, we categorized each publication into one of the seven disciplines based on the discipline of the publishing journal: neuroscience, psychology, biology, philosophy, computer science, physics, and others. To simplify the comparison across theories, we further grouped neuroscience and psychology as the “behavioral-sciences discipline” and the rest as “others” (Figure 2). We found that the proportions of publications in behavioral-sciences journals varied significantly across theories ($\chi^2(4) = 901.19$, $p < 1e-15$, Cramer’s $V = .192$). Among all theories, LRT has the highest proportion of publications in behavioral sciences (69.84%), followed by GWT (50.13%), then by HOT (46.43%), then by IIT (43.03%), and finally by QT (26.81%). We did 10 post-hoc pairwise comparisons among the theories, and found that all differences were significant ($Z > 5.42$, Bonferroni-correct $p < 1e-6$), except between GWT and HOT ($Z = 2.55$, Bonferroni-correct $p = .11$) and between HOT and IIT ($Z = 2.32$, Bonferroni-correct $p = .20$).

This pattern of results was found to be similar to that for the citation data. We performed a similar analysis on disciplines of citation sources by categorizing the disciplines of the journals in which the citing publications for each theory were published. Among all theories, LRT has the highest proportion of citations from behavioral sciences (62.56%), followed by GWT (57.93%), then by HOT (50.10%), then by IIT (45.39%), and finally by QT (27.39%). These proportions significantly varied across theories ($\chi^2(4) = 8157.8$, $p <$

1e-15, Cramer's $V = .230$). Post-hoc pairwise comparisons revealed that all these proportions significantly differed from one another (all Z values > 11 , all Bonferroni-corrected p values $< 1e-28$).

2.2 Average metrics on citation counts

Figure 3. Average metrics on citation counts across theories. LRT and QT were, respectively, the most and least cited theory across multiple metrics. Left: Average impact factors of the publishing journals across theories, with LRT being the highest and QT being the lowest. Middle: Average number of citations for each publication across theories. Publications for LRT and GWT had more citations on average than the other three theories, with QT having the least number of per-publication citations on average. Right: Proportion of publications that had at least one citation across theories, with LRT being the highest (> 90%) and QT being the lowest (< 80%). All error bars represent 95% confidence intervals. Only significant pairwise differences are marked above the bars (Bonferroni-corrected).

To compare publishing outlets across the theories, we extracted the impact factor of the publishing journal for each publication under each theory. We used the impact factor in

2018 (citations in 2018 to items published in the journal in 2016 and 2017 divided by the total number of citable items published in the journal in 2016 and 2017; Clarivate, 2020), which was the latest available metric at the time of data collection.

The left panel of Figure 3 shows the results. A one-way ANOVA revealed that average impact factors differed across theories ($F(4, 5630) = 64.72, p < 2e-16, \eta^2 = .044$). The mean impact factor was the highest for LRT (4.74), followed by GWT (4.17), then by HOT (3.76), then by IIT (3.69), and finally by QT (1.49). Except for the differences between GWT and HOT, between GWT and IIT, and between HOT and IIT (Bonferroni-corrected $p = .61, .11$, and 1 , respectively), all pairwise differences among the theories were significant (Bonferroni-corrected p values $< .015$).

To compare the number of citations across theories, we conducted a one-way ANOVA by taking each publication as an observation, with the number of citations being the dependent variable and the associated theory as the independent variable. The middle panel of Figure 3 shows the average number of citations across publications for each theory. We found that the number of citations differed across theories ($F(4, 8025) = 22.61, p < 2e-16, \eta^2 = .011$), with LRT having the largest number of citations ($M = 42.66$), closely followed by GWT ($M = 42.21$), then by IIT ($M = 26.90$), then by HOT ($M = 19.99$), and finally by QT ($M = 13.07$). Bonferroni-corrected pairwise comparisons revealed that, except between GWT and LRT (uncorrected $p = .90$), between HOT and IIT (Bonferroni-corrected $p = .97$), and between HOT and QT (uncorrected $p = .11$), the

mean number of citations differed across all pairs of theories (all Bonferroni-corrected p values < .004).

The right panel of Figure 3 shows the proportion of publications having at least one citation. We found that this “cited proportion” differed significantly across theories ($\chi^2(4) = 188.67$, $p < 1e-15$, Cramer’s $V = .153$), with LRT being highest (93.55%), followed by GWT (84.74%), then by IIT (82.44%), then by HOT (81.27%), and finally by QT (76.19%). Bonferroni-corrected pairwise comparisons revealed that, except between GWT and HOT (Bonferroni-corrected $p = .12$), between GWT and IIT (Bonferroni-corrected $p = .56$), and between HOT and IIT (uncorrected $p = .45$), the proportion of cited publications significantly differed across theories (Bonferroni-corrected p values < .03).

2.3 Nature of citations

Figure 4. Nature of citations measured as test proportions of citing publications across theories. Each citation was first categorized by scite.ai (<https://scite.ai>) as supporting the theory (light portion), contradicting the theory (dark portion), or just mentioning the theory (not shown in chart). Test proportion was defined as the sum of the supporting and contradicting proportions. Test proportion was the highest for citations of LRT and lowest for QT. The overall pattern across the theories is consistent with the patterns observed in other citation metrics (e.g., behavioral-science proportions in publications and citations and average journal impact factor; see Figure 2 and left panel of Figures 3, respectively). Each error bar represents the 95% confidence intervals for the test

proportion for each theory. Only significant pairwise differences are marked above the bars (Bonferroni-corrected).

To evaluate the quality of the citations for each theory, we analyzed the contexts of citations. Instead of merely “mentioning” the theory, meaningful citations should provide empirical evidence “supporting” or “contradicting” the theory. We used scite.ai to analyze the contexts of all citations for each theory and found that over 90% of the citations for each theory were merely “mentioning” the theory. In the other two citing contexts, the citing articles provided evidence either “supporting” or “contradicting” the theory. We grouped these two under the category of “Test” context, so that each citing article was either mentioning or had tested each of the theories.

Figure 4 shows the proportion of citing publications in the test context (with the breakdown for supporting and contradicting contexts represented by shades). This test proportion was found to be significantly different across theories ($\chi^2(4) = 533.5$, $p < 1e-15$, Cramer’s $V = .0550$). Among the theories, LRT had the highest proportion of test citations (9.07%: Supporting = 8.26%, Contradicting = 0.81%), followed by GWT (7.59%: Supporting = 6.93%, Contradicting = 0.66%), then by HOT (6.94%: Supporting = 6.40%, Contradicting = 0.54%), then by IIT (6.51%: Supporting = 6.15%, Contradicting = 0.37%), and finally by QT (3.01%: Supporting = 2.84%, Contradicting = 0.17%). A Bonferroni-corrected pairwise comparison revealed that the proportion of test citations differed significantly across all theories (Bonferroni-corrected $p < 2e-8$), except between

GWT and HOT (Bonferroni-corrected $p = .063$) and between HOT and IIT (Bonferroni-corrected $p = .80$).

Although the majority of the test citations for all theories were supporting citations, the citations that are meaningful for the development of a theory are those that provide evidence to contradict the theory. We examined the supporting and contradicting citations among the test citations for all five theories, and found that the split between contradicting and supporting citations varied across the theories ($\chi^2(4) = 32.456$, $p < 2e-6$, Cramer's $V = .0515$). We computed the contradict-to-support ratio for each theory (as a percentage, i.e., the number of contradicting citations for every 100 supporting citations), and found that the contradict-to-support ratio was the highest for LRT (9.74%), closely followed by GWT (9.55%), and then by HOT (8.39%). The contradict-to-support ratio appeared to be lower for IIT (5.95%) and QT (5.88%) than for the other three theories. In post-hoc pairwise comparisons (with Bonferroni correction), we found that the contract-to-support ratio significantly differed between LRT and IIT ($p < 3e-6$) and between GWT and IIT ($p < 2e-5$).

3. Impact in social media

3.1 Tweet counts across all publications

*: $p < .05$
**: $p < .01$
***: $p < .001$
(with Bonferroni correction)

Figure 5. Metrics from Twitter activities across the theories. Across all metrics, IIT has the highest Twitter activity across metrics. Top: Tweet counts per year from 2011 to 2019 across theories. IIT has the fastest growth and has been leading in tweet counts (except for 2018) over the majority of the survey period. Bottom-left: Average number of tweets each publication has received at the time of the survey for each theory. Bottom-right: Proportion of publications with at least one tweet mention. The pattern is similar between the two metrics, with IIT being the highest in both metrics. Each error bar represents the 95% confidence intervals for the test proportion for each theory. Only significant pairwise differences are marked above the bars (Bonferroni-corrected).

The top panel in Figure 5 shows the yearly trend in Tweet counts for all theories from 2011 to 2019. Since 2013, IIT has been the most-tweeted theory, followed by GWT, LRT, HOT, and, finally, by QT. This rank order remained unchanged until 2018, in which both IIT and GWT had a drop and HOT had a sharp increase in tweet count, making HOT the most-tweeted theory for that year. In 2019, the tweet counts decreased for both HOT and LRT but increased for IIT and GWT, restoring the previous rank order since 2013, except that HOT has overtaken LRT to become the third-most tweeted theory.

To compare the number of tweets across theories, we conducted a one-way ANOVA similar to the one we did for citation statistics. The bottom-left panel of Figure 5 shows the average number of tweets for each theory.

One publication under HOT has received 2920 tweets at the time of our survey, which corresponds to a Z score of 24.29 among publications within HOT, and a Z score of 48.24 among publications across all theories. Because that publication alone would raise the mean tweet count by 4.43 tweets for HOT (from 14.21 to 18.64), and the standard deviations by 82 tweets (from about 37.67 to 119.45), we decided to remove it from the analysis of average tweet counts over publications in order to obtain more reliable estimates.

We found that average tweet counts differed across theories ($F(4,4940) = 8.04$, $p = 2e-6$, $\eta^2 = .0065$), with IIT having the largest average number of tweets ($M = 18.45$), followed by HOT ($M = 14.21$), then by GWT ($M = 13.73$), then by LRT ($M = 10.40$), and finally by QT ($M = 8.05$). Bonferroni-corrected pairwise comparisons revealed that the

average number of tweets was higher for IIT than both for LRT (Bonferroni-corrected $p = .0088$) and for QT (Bonferroni-corrected $p = .0047$). The difference in average tweet counts between all other pairwise comparisons were not statistically significant (Bonferroni-corrected $p > .06$).

As in the analysis on citations, we computed the proportion of publications that had received at least one tweet mention. The bottom-right panel of Figure 5 shows the results. We found that IIT has the highest proportion publications having at least one tweet mention (74.43%), followed by HOT (71.54%), then by GWT (68.62%), then by LRT (62.18%), and finally by QT (61.28%). Bonferroni-corrected pairwise comparisons revealed that, except between GWT and HOT (uncorrected $p = .21$), between HOT and IIT (uncorrected $p = .19$), and between LRT and QT (uncorrected $p = .75$), the proportion of tweeted publications significantly differed across theories (Bonferroni-corrected p values $< .023$).

Our findings suggest that, across multiple metrics, Twitter activity was the highest for IIT than for other theories.

3.2 Sources of tweets by profession

Figure 6. Tweet counts by profession of the tweet source across theories.

Twitter has also been used by scientists and related professional communities, the activity of Twitter could reflect some level of academic impact. To clearly delineate the two, we attempted to examine the source of tweets across theories.

In order to extract the sources of tweets across theories, we needed to specify a set of target papers to obtain the Twitter data. We chose to focus on the three core papers for each theory (i.e., among all the papers that originally proposed/described the respective

theory, the three papers that have the largest number of citations at the time of the survey). We believe the core papers are representative of the papers for each theory, because the all-papers and core-papers data produced similar patterns of results across a number of metrics, such as total citation counts and the proportions of citations that tested (supported or contradicted) the theories (see Figures S2 and S3 in Supplementary Materials). Twitter users who mentioned the core papers were grouped by Twitter into five professions, namely, members of the public, scientists, science communicators (e.g., journalists, bloggers, editors), practitioners (e.g., doctors and other healthcare professionals), and unknown.

Figure 6 shows the tweet counts for the core papers of each theory by profession. We found that, at the time of our survey, the numbers of tweet mentions of the core papers for QT (258 tweets) and IIT (199 tweets) were much larger than those for HOT (16 tweets), LRT (7 tweets), and GWT (5 tweets). For all theories, most of the tweet mentions came from members of the public. The proportion of tweet mentions from the public was higher for QT (219 out of 258; 84.88%) and IIT (154 out of 199; 77.39%) than that for GWT (3 out of 5; 60%), HOT (9 out of 16; 56.25%), and LRT (3 out of 7; 42.86%).

Because there was a huge discrepancy in the numbers of tweet mentions between GWT, LRT, and HOT (with only 5, 7, and 16 tweets, respectively) and the two more well-tweeted theories (QT and IIT, with 258 and 199 tweets respectively), the statistical tests for pairwise comparisons across all five theories would be unreliable. In order to address this issue, we grouped the data for GWT, LRT, and HOT as a group of “Others”,

with 15 out of 28 tweets (53.57%) coming from the public. Bonferroni-corrected pairwise comparisons revealed that the proportion of “public tweets” was higher for IIT (77.39%; Bonferroni-corrected $p = .04$) and for QT (84.88%; Bonferroni-corrected $p = .0004$) than for the other three theories (53.57%), though we did not find any difference in the public-tweet proportion between IIT and QT (uncorrected $p = .054$).

4. Bibliometric similarities among theories

Figure 7. Bibliometric dissimilarities (pairwise) among the five theories for consciousness. Left: Pairwise dissimilarities across all 14 bibliometric dimensions. Right: 2D-scaled visualization of the bibliometric “distances” among the theories. IIT, LRT, and QT appear to differ the most among all theories, while GWT, HOT, and LRT appear to form a cluster among themselves.

To evaluate how the theories differ across the bibliometric measures, we computed the pairwise multivariate distances as a measure of dissimilarities among the theories. We included 14 dimensions. The 11 dimensions for Academia are the 4 growth rates (publication, citations, citations for core paper, citation per publication; Figure 1), the 2 proportions (in log-odds) of behavioral-science outlets (publication and citations; Figure 2), the 3 metrics on citation counts (average Journal Impact Factor, average citation per publication, and log-odds of proportions of cited publications), the tested proportions in citations (bars in Figure 4) and the log contradict-to-support ratios (dark-to-light bars in

Figure 4). The 3 dimensions for Social media are average tweet counts per publication (Figure 5, bottom-left), log-odds of tweeted publications (Figure 5, bottom-right), and log-odds of tweet of core papers coming from public (Figure 6, light bar over whole bar).

To allow each dimension to be weighted equally, we standardized the scores for the theories in each dimension by computing the Z scores, i.e., for each dimension, we subtracted the mean across theories from the raw scores and divided the difference by the standard deviation across theories. Then we computed the Euclidean distance $d_{i,j}$ between two theories i and j as follows

$$d_{i,j} = \left[\sum_{k=1}^K (z_{k,i} - z_{k,j})^2 \right]^{1/2}$$

, where $z_{k,i}$ denotes the standardized score in dimension k for theory i , and $K = 14$ denotes the total number of dimensions.

The mean and standard deviation of the 10 pairwise distances were, respectively, 5.09 and 1.46. Among all distances, the shortest was between GWT and HOT (2.90), suggesting that they were the most similar theories across the bibliometric measures.

The longest distance was between LRT and QT (7.51), suggesting that they differed the most among all theory pairs in bibliometric measures. In particular, among GWT, HOT, and LRT (average pairwise distance = 3.52) appeared to form a theory cluster, while IIT (closest to HOT with pairwise distance = 3.61; average pairwise distance with other theories = 5.33) and QT (closest to HOT with pairwise distance = 5.45; average pairwise distance with other theories = 6.28) appeared to be isolated from their theories.

This suggests that the theories of consciousness can be categorized into three groups based on their impact in academia and social media, namely, IIT, QT, and the GWT-HOT-LRT cluster.

DISCUSSION

Overall, we found that the impact of the five theories for consciousness varied across platforms. Specifically, within academia, we found that GWT and LRT have greater impact than other theories do across multiple metrics, including citation counts per publication, average impact factor across publishing journals, and the relative frequency of being tested in their respective citations. IIT is the fastest-growing theory in terms of publication and citation counts, but the proportion of citations that empirically tested it was relatively low and is not commensurate with its growth. In social media, IIT has greater impact than other theories do, leading in absolute tweet counts and proportions of tweeted publications. QT appears to have the least impact across metrics in academia, despite its exceptionally-high number of tweets on social media. While both IIT's and QT's core papers appear to have a much higher impact on social media than other theories do, most of their Twitter mentions came from the public. We discuss below some of the key findings from the present study.

Our analysis on growth rates in publication and citation counts reveals that a theory could grow fast on raw counts, but slow on average citation counts per publication. Specifically, in terms of growth rate in both publication and citation counts (Figure 1, top two panels), IIT is the fastest-growing theory (consistent with Figure 2 in Jake Hanson's blog), while QT has the slowest growth rates among the theories. However, for the growth in average citation counts per publication, (Figure 1, bottom panel), IIT is the

slowest-growing theory. This suggests that its growth in citation has not been catching up with its growth in publication. This phenomenon appears to be specific to IIT, because the second-fastest-growing theory in raw publication and citation counts (HOT) is also the second-fastest-growing theory in per-publication citation counts.

In addition to growth rates, our analysis of citations shows that theories differ in their levels of reception in academia. LRT is the top in terms of all average-citation metrics (Figure 3), with a close-tie with GWT in per-publication citation counts (middle panel, Figure 3). Within its own citations, LRT is also the most frequently-tested theory (Figure 4). Across all the citation metrics we have analyzed, QT has been consistently ranked the lowest. In addition, the overall pattern in the proportions of citations from behavioral-science disciplines is similar to that in the test proportions of citations. In particular, despite its fastest growth in publication and citation counts over the years, IIT is the second-lowest in terms of test rate within its citing publications, significantly lower than those for GWT and LRT. This implies that the fast-growing citations that IIT has attracted over the years did not offer as much empirical evidence to support or contradict the theory as the citations for theories like GWT and LRT did.

Why were citations of IIT and QT “mentioning” them instead of “testing” them? One possible reason is that a large proportion of citations for the two theories did not come from behavioral-science disciplines such as neuroscience and psychology, in which theories of consciousness have been most actively studied and empirically tested. The

patterns of results for test rates (Figure 4) and proportions of citations coming from behavioral-science disciplines (Figure 2) are very similar, suggesting a possible relationship between these two metrics. Specifically, among all theories, LRT has the highest proportion of citations coming from neuroscience and psychology and, at the same time, is the most frequently-tested among its citing publications. The same logic, despite being in the other direction, could be applied to understand the relationship between IIT's and QT's low proportions of citations from neuroscience and psychology and their low proportions of test citations. Specifically, for IIT, the low test rate may be related to the idea that IIT, as a theory for consciousness, is inherently difficult, if possible at all, to falsify (Hanson & Sara, 2020).

In social media, we found that the Twitter activity for IIT and QT was higher, across multiple metrics, than that for GWT, HOT, and LRT. For IIT, this could be related to IIT's fast growth in raw publication and citation counts. In particular, when we focused on the core papers of each theory, IIT and QT had many more tweets (around 200 each) than other theories did (only 28 even with all theories combined). Among their many twitter mentions, the majority (around 80%) came from the public instead of from the science of professional communities. For the other theories, only around half of the twitter mentions came from the public.

This finding is consistent with the results in a survey conducted by Michel et al. (2018), which showed that IIT was considered the most promising among all theories only

among non-experts. One possibility is already mentioned in Michel et al. (2018): the perception that IIT is promising may come from its promotion among the public but not as much among researchers in the field. Our findings seem to be compatible with this “overpromotion” hypothesis.

Another possible explanation for the high Twitter activity specifically on IIT’s core papers is timing: two of IIT’s core papers were published when Twitter started becoming popular (2008, 2014), while some core papers of other theories had been published before Twitter was available (e.g., GWT: 1988 - 2001; HOT: 1986 & 2005; LRT: 2000 & 2003). However, this explanation does not apply to QT’s high Twitter activity, because two of QT’s core papers had also been published before Twitter’s time (1996 & 1998).

Overall, the above findings can be summarized in the bibliometric similarities shown in Figure 7, in which the theories of consciousness can be grouped into three types. GWT, HOT, and LRT form a theory cluster, marked by their moderate growth rates in academic metrics, in-depth impact in relevant disciplines, and relatively low impact (especially on the public) in social media. IIT differs from this theory cluster in that it has high growth rate in counts but relatively shallow impact in academia, especially on the relevance of disciplines and extent of being empirically tested, plus high impact in social media, especially with tweets from the public. QT differs from the rest in that it has the lowest impact almost across all metrics in academia but high impact in social media.

We acknowledge that there are both methodological and statistical limitations in the present study. We attempt to address them in the discussion below.

One could doubt the reliability of the mention/support/contradict classification by scite.ai, because the machine-learning algorithm behind scite.ai was built only on a subset (around 16 million full-text articles that are NOT behind paywalls) of all research articles (e.g., over 53 million articles on Web of Science and 112 million articles registered on Crossref (Khamisi, 2020)). Also, some citations might have been misclassified because the accuracy is not 100%. However, we ran the same algorithm with the same settings across all five theories. Assuming that the algorithm has no specific bias towards any particular theories, we believe it is still legitimate to compare the results across theories. Practically, this algorithm allows researchers to perform in-depth analysis on citation contexts over tens of thousands of articles only with a few clicks, which is helpful especially for macroscopically evaluating the quality of citations for a theory. To provide more reliable analyses on citation quality, future studies could employ methods such as a systematic review to analyze the citing contexts for the theories.

Furthermore, in the present study, statistical significance could have been boosted by the large sample size, especially for the per-publication analyses. This may lead to the detection of effects with very small effect sizes. Because of this, instead of interpreting only one small effect, we attempted to base our discussion and conclusions on the

overall pattern across multiple metrics, which has provided us with the big picture. For example, for the per-publication citation analyses, LRT was consistently top for journal impact factors, average citation counts, and cited proportions; for Twitter activity, IIT was consistently the top for per-year tweet counts, per-publication tweet counts, and tweeted proportion.

On the other hand, for some metrics (e.g., sources of tweets for the core papers for GWT, HOT, and LRT), the number of observations could be insufficient for us to perform reliable statistical analyses. We acknowledge this and we have avoided making strong conclusions based on those findings.

In conclusion, we found that theories of consciousness varied substantially in their impact across both in academia and on social media. Specifically, in terms of raw numbers, IIT was the fastest-growing theory for consciousness in both academia and social media. However, unlike other theories of consciousness such as GWT and LRT, IIT's growth in quantity was not commensurate with the nature of its impact. This appeared to be a phenomenon specific to IIT because such discrepancy was not observed for other theories with similar growth rate in raw counts. Our findings suggest that this apparent lack of impact quality could be explained by media promotion to the general public, over its academic impact to the relevant expert groups.

REFERENCES

- Baars, Bernard J. (1988). *A cognitive theory of consciousness*. Cambridge [England] ; New York : Cambridge University Press,
<http://www.loc.gov/catdir/toc/cam032/87020923.html>
- Botting, N., Dipper, L., & Hilari, K. (2017). The effect of social media promotion on academic article uptake. *Journal of the Association for Information Science and Technology*, 68(3), 795–800. <https://doi.org/https://doi.org/10.1002/asi.23704>.
- Brown, R., Lau, H., & LeDoux, J. E. (2019). Understanding the Higher-Order Approach to Consciousness. *Trends in Cognitive Sciences*, 23(9), 754–768.
<https://doi.org/10.1016/j.tics.2019.06.009>.
- Clarivate. (2020). The Clarivate Analytics Impact Factor.
<https://clarivate.com/webofsciencegroup/essays/impact-factor/>
- Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition*, 79(1), 1–37.
[https://doi.org/10.1016/S0010-0277\(00\)00123-2](https://doi.org/10.1016/S0010-0277(00)00123-2).
- Dehaene, S., Kerszberg, M., & Changeux, J.-P. (1998). A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the National Academy of Sciences*, 95(24), 14529–14534. <https://doi.org/10.1073/pnas.95.24.14529>.
- Doerig, A., Schurger, A., Hess, K., & Herzog, M. H. (2019). The unfolding argument: Why IIT and other causal structure theories cannot explain consciousness. *Consciousness and Cognition*, 72, 49–59.
<https://doi.org/10.1016/j.concog.2019.04.002>.

- Ekert, A., Jozsa, R., Penrose, R., & Stuart, H. (1998). Quantum computation in brain microtubules? The Penrose–Hameroff ‘Orch OR’ model of consciousness. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 356(1743), 1869–1896.
<https://doi.org/10.1098/rsta.1998.0254>.
- Hameroff, S., & Penrose, R. (1996). Orchestrated reduction of quantum coherence in brain microtubules: A model for consciousness. *Mathematics and Computers in Simulation*, 40(3), 453–480. [https://doi.org/10.1016/0378-4754\(96\)80476-9](https://doi.org/10.1016/0378-4754(96)80476-9).
- Hameroff, S., & Penrose, R. (2014). Consciousness in the universe: A review of the ‘Orch OR’ theory. *Physics of Life Reviews*, 11(1), 39–78.
<https://doi.org/10.1016/j.plrev.2013.08.002>.
- Hanson, J.R. (2020, June 24). My Experience with Integrated Information Theory (IIT). Jake Hanson’s Blog.
<http://jakerhanson.weebly.com/3/post/2020/06/my-graduate-experience-with-integrated-information-theory-iit.html>.
- Hanson, Jake R., & Walker, S. I. (2020). Formalizing Falsification for Theories of Consciousness Across Computational Hierarchies. *ArXiv:2006.07390* [Cs, Math].
<http://arxiv.org/abs/2006.07390>.
- Khamsi, R. (2020). Coronavirus in context: Scite.ai tracks positive and negative citations for COVID-19 literature. *Nature*. Epub ahead of print. doi: 10.1038/d41586-020-01324-6.
<https://www.nature.com/articles/d41586-020-01324-6>.

Kleiner, J. (2020). Brain states matter. A reply to the unfolding argument.

Consciousness and Cognition, 85, 102981.

<https://doi.org/10.1016/j.concog.2020.102981>.

Lamme, V. A. F. (2003). Why visual attention and awareness are different. *Trends in*

Cognitive Sciences, 7(1), 12–18. [https://doi.org/10.1016/S1364-6613\(02\)00013-X](https://doi.org/10.1016/S1364-6613(02)00013-X).

Lamme, V. A. F. (2006). Towards a true neural stance on consciousness. *Trends in*

Cognitive Sciences, 10(11), 494–501. <https://doi.org/10.1016/j.tics.2006.09.001>.

Lamme, V. A., & Roelfsema, P. R. (2000). The distinct modes of vision offered by

feedforward and recurrent processing. *Trends in Neurosciences*, 23(11), 571–579.

[https://doi.org/10.1016/s0166-2236\(00\)01657-x](https://doi.org/10.1016/s0166-2236(00)01657-x) .

Lau, H., & Rosenthal, D. (2011). Empirical support for higher-order theories of conscious

awareness. *Trends in Cognitive Sciences*, 15(8), 365–373.

<https://doi.org/10.1016/j.tics.2011.05.009>.

Michel, M., Beck, D., Block, N. et al. Opportunities and challenges for a maturing

science of consciousness. *Nat Hum Behav* 3, 104–107 (2019).

<https://doi.org/10.1038/s41562-019-0531-8>.

Michel, M., Fleming, S. M., Lau, H., Lee, A. L., Martinez-Conde, S., Passingham, R. E.,

. . . Liu, K. (2018). An informal internet survey on the current state of consciousness science. *Frontiers in Psychology*, 9, 2134.

Oizumi, M., Albantakis, L., & Tononi, G. (2014). From the Phenomenology to the

Mechanisms of Consciousness: Integrated Information Theory 3.0. *PLOS*

Computational Biology, 10(5), e1003588.

<https://doi.org/10.1371/journal.pcbi.1003588>,

Rosenthal, D. M. (1986). Two Concepts of Consciousness. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 49(3), 329–359.

<https://www.jstor.org/stable/4319832>.

Rosenthal, D. M. (2005). *Consciousness and Mind*. Oxford University Press.

Tononi, G. (2004). An information integration theory of consciousness. *BMC Neuroscience*, 5(1), 42. <https://doi.org/10.1186/1471-2202-5-42>.

Tononi, G. (2008). Consciousness as integrated information: a provisional manifesto. *The Biological Bulletin*, 215(3), 216–242. <https://doi.org/10.2307/25470707>.

Tsuchiya, N., Andriellon, T., & Haun, A. (2019). A reply to “the unfolding argument”: Beyond functionalism/behaviorism and towards a truer science of causal structural theories of consciousness. *PsyArXiv*. <https://doi.org/10.31234/osf.io/a2ms9>.

SUPPLEMENTARY MATERIALS

Core papers

The three publications with the highest number of citations on Scopus within each theory were taken as the “core papers” for that theory. Below is the list of core papers for each theory.

GWT

1. Baars, Bernard J. (1988). *A cognitive theory of consciousness*. Cambridge [England] ; New York : Cambridge University Press, <http://www.loc.gov/catdir/toc/cam032/87020923.html> [1750 citations]
2. Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition*, 79(1), 1–37. [https://doi.org/10.1016/S0010-0277\(00\)00123-2](https://doi.org/10.1016/S0010-0277(00)00123-2) [1277 citations]
3. Dehaene, S., Kerszberg, M., & Changeux, J.-P. (1998). A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the National Academy of Sciences*, 95(24), 14529–14534. <https://doi.org/10.1073/pnas.95.24.14529> [636 citations]

HOT

1. Rosenthal, D. M. (2005). *Consciousness and Mind*. Oxford University Press. [399 citations]
2. Rosenthal, D. M. (1986). Two Concepts of Consciousness. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 49(3), 329–359. <https://www.jstor.org/stable/4319832> [372 citations]

3. Lau, H., & Rosenthal, D. (2011). Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Sciences*, 15(8), 365–373.
<https://doi.org/10.1016/j.tics.2011.05.009> [250 citations]

IIT

1. Tononi, G. (2004). An information integration theory of consciousness. *BMC Neuroscience*, 5(1), 42. <https://doi.org/10.1186/1471-2202-5-42> [746 citations]
2. Tononi, G. (2008). Consciousness as integrated information: a provisional manifesto. *The Biological Bulletin*, 215(3), 216–242. <https://doi.org/10.2307/25470707> [575 citations]
3. Oizumi, M., Albantakis, L., & Tononi, G. (2014). From the Phenomenology to the Mechanisms of Consciousness: Integrated Information Theory 3.0. *PLOS Computational Biology*, 10(5), e1003588.
<https://doi.org/10.1371/journal.pcbi.1003588> [275 citations]

LRT

1. Lamme, V. A., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences*, 23(11), 571–579.
[https://doi.org/10.1016/s0166-2236\(00\)01657-x](https://doi.org/10.1016/s0166-2236(00)01657-x) [1357 citations]
2. Lamme, V. A. F. (2003). Why visual attention and awareness are different. *Trends in Cognitive Sciences*, 7(1), 12–18. [https://doi.org/10.1016/S1364-6613\(02\)00013-X](https://doi.org/10.1016/S1364-6613(02)00013-X) [647 citations]

3. Lamme, V. A. F. (2006). Towards a true neural stance on consciousness. *Trends in Cognitive Sciences*, 10(11), 494–501. <https://doi.org/10.1016/j.tics.2006.09.001> [449 citations]

QT

1. Hameroff, S., & Penrose, R. (2014). Consciousness in the universe: A review of the ‘Orch OR’ theory. *Physics of Life Reviews*, 11(1), 39–78.
<https://doi.org/10.1016/j.plrev.2013.08.002> [256 citations]
2. Ekert, A., Jozsa, R., Penrose, R., & Stuart, H. (1998). Quantum computation in brain microtubules? The Penrose–Hameroff ‘Orch OR’ model of consciousness. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 356(1743), 1869–1896.
<https://doi.org/10.1098/rsta.1998.0254> [189 citations]
3. Hameroff, S., & Penrose, R. (1996). Orchestrated reduction of quantum coherence in brain microtubules: A model for consciousness. *Mathematics and Computers in Simulation*, 40(3), 453–480. [https://doi.org/10.1016/0378-4754\(96\)80476-9](https://doi.org/10.1016/0378-4754(96)80476-9) [122 citations]

Breakdown by discipline of publishing outlets and citing outlets across theories.

Figure S1. Breakdown by discipline of publishing and citing outlets across theories.

Error bars are 95% confidence intervals for the proportions of “behavioral-sciences” disciplines (i.e., psychology + neuroscience) as indicated in Figure 2.

Similarities between all-papers and core-papers metrics

Figure S2. Average number of citations per paper for all papers (left) and core papers (right). The overall pattern is similar, except that the core papers for LRT had fewer citations than its average number across all papers.

Figure S3. Proportions of tested citations for each theory, plotted across all papers (left) and core papers (right). Error bars are 95% CIs, which are longer for the core papers because the proportions were estimated based on fewer citations for the core papers (around 600-3,000) than for all papers (around 10,000 to 60,000).