© *Academy of Management Journal*
2018, Vol. 61, No. 5, 1613–1625.
https://doi.org/10.5465/amj.2018.4005

FROM THE EDITORS

A BRIEF PRIMER ON DATA VISUALIZATION OPPORTUNITIES IN MANAGEMENT RESEARCH

"A picture is worth a thousand words." A picture or illustration not only helps readers to grasp the key content of a study, but also to remember it. A picture or illustration provides alternative mechanisms for communicating, akin to speaking a second language. In fact, research has long shown that individuals understand (Carney & Levin, 2002) and better remember information when it is delivered in pictures than when delivered in single words or short sentences (Shepard, 1967). Traditional data visualization methods, such as scatter plots, bar charts, histograms, line charts, and pie charts, are widely used in management research. In parallel to the rapid evolution of data science, however, new techniques to visualize quantitative and qualitative data have also been developed.

In this editorial, we briefly introduce some data visualization techniques used in disciplines other than management that could be particularly beneficial for management research. In the interest of space, we focus on techniques that are primarily conducted on raw data prior to formal statistical analyses. We first provide a selective overview of techniques used to visualize *qualitative data*, including word clouds, word trees, social graphs, and history flows. We then do the same for techniques used to visualize *quantitative data*, where we touch on multidimensional scaling plots, funnel plots, maps, bubble plots, dynamic plots, and tree maps. We also discuss different software packages that are useful for data visualization, and provide a few suggestions on how to make data visualization aesthetically appealing. For each technique, we offer a brief description, leaving it to researchers to explore and digest additional assumptions and caveats.[1]

VISUALIZATION OF QUALITATIVE DATA

Roughly 20% of *Academy of Management Journal* submissions are qualitative papers. Reviewers frequently note that authors of qualitative papers struggle to find informative and appealing ways to communicate and showcase their results. From the author's perspective, finding a compelling way in which qualitative data can be communicated to one's readers can also help to establish trust in the results. Four techniques with which to visualize qualitative data are discussed—and illustrated, Figures 1a and 1b, 2a and 2b, and 3—in the sections that follow.

Word Clouds

A "word cloud" is a method to visualize text in which the more frequently used words are highlighted by occupying greater prominence in the representation (McNaught & Lam, 2010). A word cloud depicts the frequencies of different words in a given text so that words that appear more frequently are larger in the cloud than those that appear less frequently. Hence, word clouds enable researchers to provide an overview of the main topics and themes in the text and identify systematic patterns—elements that are core to the work of many qualitative researchers. Comparisons can also be made across word clouds generated from different texts to highlight their thematic similarity and differences.

Word clouds have been used to offer insights from large amounts of text data. For example, Colquitt (2013) used text from the introduction sections of 174 *Academy of Management Journal* articles to create a word cloud that illustrated the relative prominence of the various research themes represented within these articles. Evanschitzky and Goergen (2018) used a word cloud to analyze the author-provided keywords of all articles published in the *British Journal of Management* between 2000 and 2015 to show that the most frequent supplied keywords were "performance," "management," "organization," "knowledge," and "strategy." Word clouds have also been applied in research areas that involve text or content analysis, such as ideology (Gentzkow, Shapiro, & Taddy, 2016) and culture (Schmiedel, Müller, & vom Brocke, 2018). Nonetheless,

We wish to acknowledge the excellent and integral research support of Tengjian Zou.

[1] We also encourage readers to look up other editorials on this topic such as Greve (2017).

a search for "word cloud" in the top management journals reveals that they are not typically presented in qualitative articles—most likely because researchers may prefer to use this technique for their own preliminary understanding in terms of what topics and themes appear in the text, and to obtain guidance on those that warrant further investigation.

Caveats. It is worth noting that raw text may be unformatted and have redundant information, which could add noise to the final result. For this reason, it is often necessary to preprocess the text before generating a word cloud. Specifically, preprocessing includes concatenating information, lowercasing the text, tokenizing the words, removing words (e.g., stop words and numbers), and stemming the words (Arts, Cassiman, & Gomez, 2018).

WordCloud, an open-source package in Python developed by Andreas Mueller,[2] can be used to generate word clouds. This package provides researchers with the flexibility to configure the size of the word cloud, number of words, background color, word font, and other aesthetic features.

We provide two examples of this visualization using text data from the commencement speeches for Stanford University, for 2016 (Burns, 2016; Figure 1a) and for 2017 (Cuéllar, 2017; Figure 1b).[3] Very briefly, the word cloud in Figure 1a immediately suggests that the main topic of the 2016 commencement speech was history (including links to American, national, and Lincoln), whereas the word cloud in Figure 1b suggests that the 2017 commencement speech revolved around the topics of the world, life, and work.

FIGURE 1A
Word Cloud Plotted Using the Stanford University 2016 Commencement Speech[a]

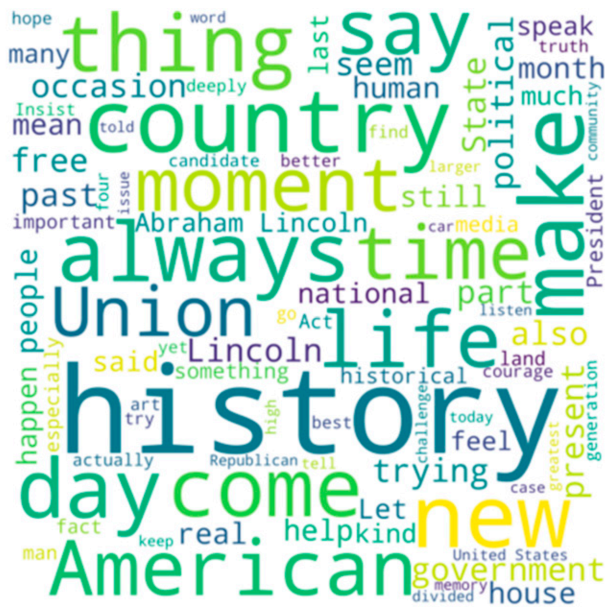

[a] Burns (2016).

FIGURE 1B
Word Cloud Plotted Using the Stanford University 2017 Commencement Speech[b]

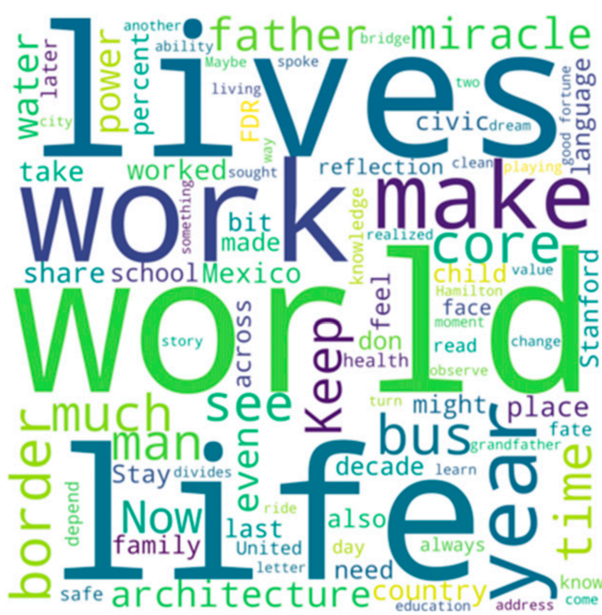

[b] Cuéllar (2017).

Word Trees

The "word tree" visualization technique was invented by Wattenberg and Viégas (2008), and offers a new information-retrieval technique for use with text data. A word tree places a tree structure onto the words or phrases that follow a particular word or phrase, and then uses that structure to arrange those words or phrases spatially. Similar to a word cloud, words or phrases that appear more frequently are larger in the graph. In addition, a word tree has a visual tree structure that makes it easy to spot repetition in the contextual words that follow a word or a phrase.

For example, Mitra and Gilbert (2014) studied what phrases positively and negatively predicted successful crowdfunding projects on Kickstarter. To explain why some phrases have positive effects and others have negative effects, they used word trees (see Figure 6 in

[2] https://amueller.github.io/word_cloud/

[3] The code for the figures we use as examples of the different visualization techniques can be found at https://github.com/tengjian-zou/data-visualization-in-python.

Mitra & Gilbert, 2014) to show that the positive phrase "pledgers will" is often followed by "receive," which conveys the information that one would receive gifts or other benefits after funding the project. In comparison, the negative phrase "even a dollar" is often followed by words such as "short," "will," "can," and "helps," which might be interpreted as desperation or groveling for money and therefore is perhaps less appealing.

Caveats. As with the case of word clouds, this technique pulls from raw text and thus it is important to ensure that the raw text does not have redundant information, which could add noise to the final result. To create word tree graphs, researchers can directly use the online platform created by Jason Davies.[4]

In Figures 2a and 2b, we provide two examples of this visualization technique using text of Apple's iPhone 8 product description and text of the poem *We Two, How Long We Were Fool'd* by Walt Whitman. Figure 2a suggests that Apple emphasizes the camera as a notable feature of the iPhone 8, and also that being the "world's…" and "most…" were important themes in their description. Figure 2b reveals the repeated structures in Whitman's poem around themes that are consistent with the poem's title and message.

Social Graphs

A "social graph" is a diagram that illustrates interconnections among people, groups, or organizations in a social network. In social graphs, each person, group, or organization (typically referred to as an "actor") is represented by a dot called a "node," and the relationships among the people, groups, or organizations are represented by lines called "edges." Social graphs can be used to display both qualitative and quantitative data. Specifically, while the width of edge and size of node are often used to depict quantitative data, the color of nodes and edges, and labels for nodes, can be used to display qualitative data. Although social graphs are used more frequently in studies with a quantitative design (including text and sentiment analysis, more recently),[5] we believe that greater use of these graphs could be particularly beneficial for visualizing qualitative data, or for visualizing a mix of qualitative and quantitative data (Jones, Maoret, Massa, & Svejenova, 2012).

Social graphs are widely employed for social network analysis. For example, Burt (2004: 352, Figure 1) created a social graph in order to compare an individual with a sparse ego network to one who has a dense ego network; he later used a social graph to illustrate direct and indirect contacts around an investment banker (Burt, 2007: 120, Figure 1). Fonti and Maoret (2016: 1772, Figure 1) created a social graph to depict members in an organizational core and those in the organizational periphery.

Caveats. When the network size is large (i.e., there are many nodes) or network density is high (i.e., there are many edges), the nodes and edges displayed on the social graph may overlap or be very close to each other, and thereby not produce an appealing or informative graph. Researchers can decrease the size of nodes, width of edges, and the luminosity of nodes and edges, or increase the size of the graph, or try different layouts (such as a circular layout, random layout, spring layout, or Kamada-Kawai layout) to provide a graph that is more helpful for their readers.

An open source package called NetworkX in Python can be used to generate social graphs, both directed and undirected (Hagberg, Swart, & Chult, 2008). With this package, one can easily add nodes and edges to a graph, and choose among different layouts (such as the ones mentioned above). Users also have the flexibility to configure the size, color, and shape of the nodes and edges. An understanding of Python programming is necessary to use NetworkX. Alternatively, scholars can use NetDraw (Borgatti, 2002), a Windows-based software with a more user-friendly graphic interface, and which does not require programming.

In Figure 3, we provide an example of this visualization using cosponsoring networks of female representatives (House) in Congress 105.[6] The red nodes in the social graph represent female representatives from the Republican Party, and the blue nodes in social graph represent female representatives from the Democratic Party. What we show in Figure 3 are directed graphs (where, if actor A co-sponsors a bill sponsored by actor B, one would see a line [edge] going from actor A to actor B), wherein the red edges represent those that originated from Republican female representatives, and the blue edges represent those that originated from Democratic female representatives. The four social graphs are generated using four different layout algorithms, as indicated in the note to the figure.

[4] https://www.jasondavies.com/wordtree/

[5] For example, Erkan and Radev (2004: 43, Figure 2) used a social graph to visualize the cosine similarity between sentences, wherein the nodes in the social graph represented the sentences and the edge was weighted to show the cosine similarity between sentences.

[6] Data collected from https://www.opensecrets.org/

FIGURE 2A
Word Tree Plotted Using Text of the Apple iPhone 8 Product Description

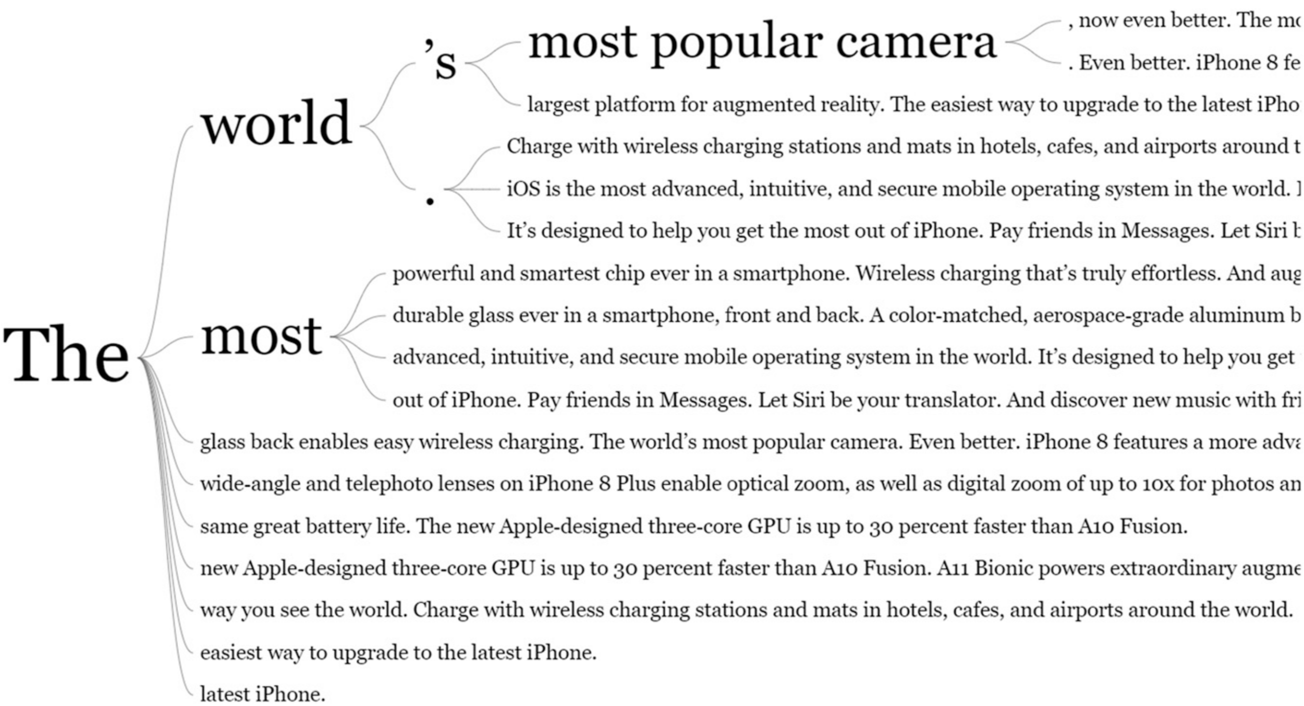

Note: The period after “The world” means that the phrase “the world” appears at the end of some sentences.

FIGURE 2B
Word Tree Plotted Using the Poem “We Two, How Long We Were Fool’d” by Walt Whitman

FIGURE 3
Social Graphs Plotted Using the Cosponsoring Networks of Female Members of the House of Representative in Congress 105

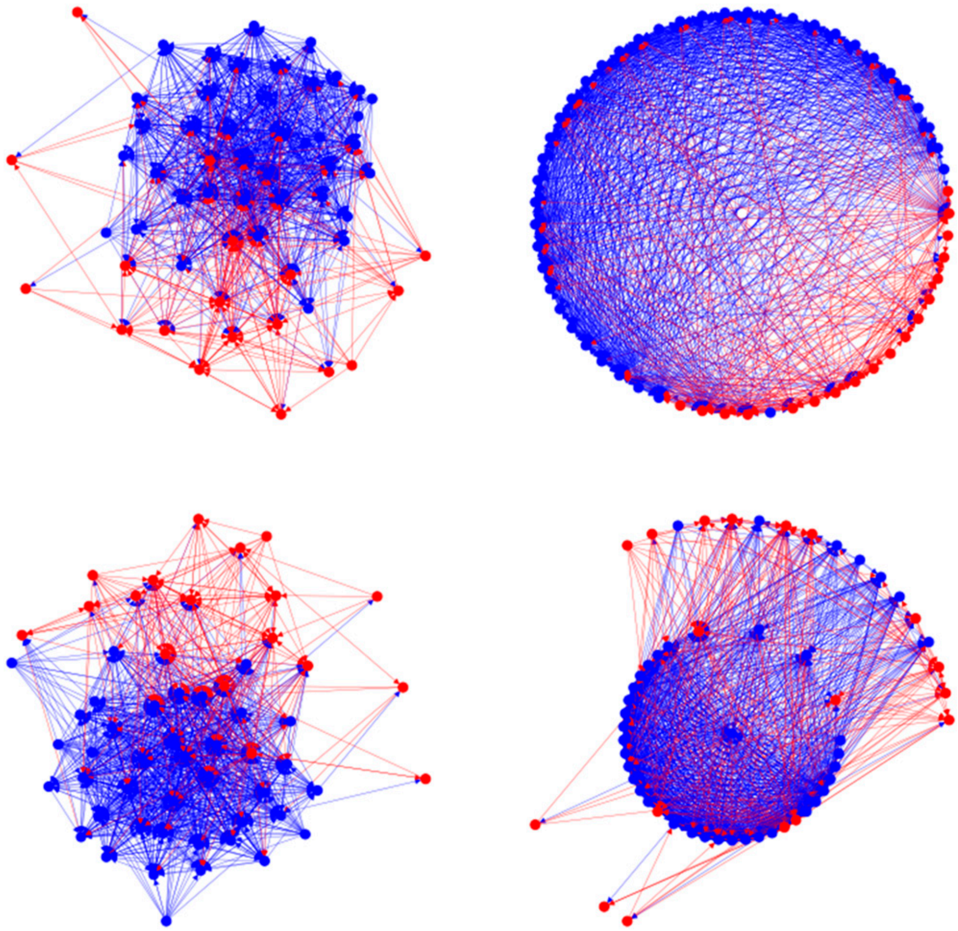

Note: The four social graphs were generated using the same data but are depicted in different layouts (upper left = Kamada–Kawai layout; upper right = circular layout; lower left = NEATO layout; lower right = TWOPI layout).

History Flows

"History flows" are a text visualization method invented by Wattenberg and Viégas (2010) in 2003, the initial purpose of which was to visualize the dynamic editing history of specific entries on Wikipedia. An example given in their book draws on the 198 versions of the Wikipedia entry for "Microsoft," where Wikipedia shows the editing history sorted by time in text for a given entry. Specifically, in a history flow graph, the *x*-axis indicates the sequence of different versions of an entry, with the first version on the left, the second version to its right, and so on. Each version has a *y*-axis, the length of which corresponds to the length of text in that version. Each person who edited the article is represented by a unique color. As a result, for each *y*-axis, the color represents a specific person, and the length represents how much text this person added. Yet, because people can add, delete, and insert text, and because of the large number of people who might have worked on the same entry, it is often difficult to grasp how people interact with one another in each version and across different versions.

History flow graphs have also been used to visualize the interactions among authors on Wikipedia. For example, Viégas, Wattenberg, and Dave (2004) drew on this technique to reveal a complex history of cooperation and conflict among authors on Wikipedia, so that it can be observed which authors worked on the same version and whether they added, deleted, or inserted text.

Based on our brief review, history flow graphs have yet to be used in management research. They have the potential, however, to enhance the visualization of histories of cooperation between individuals and firms. For example, in research that uses data from online open source software (OSS) development communities, researchers can draw on history flow graphs to visualize those who worked on the same OSS projects and how the cooperation patterns change over time. In addition, when researchers analyze online user reviews, user history flow graphs can display which users review the same product and their ratings for that product. Readers can then develop a better sense of the dynamics of product popularity or user taste.

Caveats. IBM's History Flow tool was originally used to plot history flows, but is no longer available for download from the IBM research website.[7] We call for researchers to contribute to open source code to implement this technique for future use by researchers in management, as well as those in other applicable fields of research.

VISUALIZATION OF QUANTITATIVE DATA

Management researchers frequently use illustrations of quantitative data such as scatter plots, bar, line, and pie charts, and histograms. Some novel and less frequently seen ways of visualizing quantitative data are discussed briefly in the following sections, and depicted in Figures 4, 5, 6, and 7 below.

Multidimensional Scaling Plots

The purpose of multidimensional scaling (MDS) is to provide a visual representation of proximities (i.e., similarities or distances) among a set of objects. MDS enables users to reduce dimensionalities so that the objects can be displayed on two-dimensional or three-dimensional plots. For example, Hofstede's (1980) culture measure originally had four dimensions (power distance, individualism, uncertainty avoidance, and masculinity). By using these four dimensions, we can calculate, for example, the Euclidean cultural distance (Kogut & Singh, 1988) for all pairs of countries, and plot these distances on a two-dimensional map, such that those countries that are culturally similar to each other are placed nearer to each other on the map than those that are culturally different. Robinson and Bennett (1995: 562–563, Figure 1) generated a two-dimensional scaling plot to compare the perceived differences between deviant behavior descriptions. Pentland, Hærem, and Hillison (2011: 1374, Figure 2) used two-dimensional scaling plot to compare the differences among various sequences of routines. In these three examples, culture, deviant behavior, and routines originally have multiple dimensions, but MDS reduces these dimensions to two, by showing the distances (e.g., Euclidean distance) among all pairs of relevant objects.

Caveats. To decide how many dimensions are needed for an MDS plot, researchers should check the "stress" from each dimension. Stress is an indicator of goodness of fit, indicating how accurate the MDS plot can represent the data. The smaller the stress, the better the MDS plot can represent the data (Kruskal, 1964: 3).

XLSTAT, a commercial software program that is embedded in Microsoft Excel after installation, can be used to generate MDS plots. Before producing the MDS plot, one needs to prepare an n-by-n matrix, where each element in the matrix is the similarity or distance between the corresponding row/column. For example, take the aforementioned cultural distance MDS plot: n indicates n countries, and each element in the matrix denotes the cultural distance between country i and country j. XLSTAT can directly produce the stress values for different dimensions, and users can choose the dimensions with reasonable stress values. It is convenient in XLSTAT to move and edit the label for each object in the plot, which can help to avoid issues with overlapping labels in some other software.

We provide an example of this visualization in Figure 4 using country cultural distance data. The plot includes 83 countries or areas. The cultural distance was calculated using the Kogut and Singh (1988) index (for more on this index, see Cuypers, Ertug, Heugens, Kogut, & Zou, in press), based on Hofstede's (1980) original four culture dimensions (power distance, individualism, uncertainty avoidance, and masculinity).

Funnel Plots

Meta-analytic studies occupy an important place in the field of management (Shaw & Ertug, 2017). One of the key concerns about meta-analysis is publication bias, whereby studies with stronger effects are more likely to be published, inflating the final meta-analytic results (Rothstein, Sutton, & Borenstein, 2006). To be transparent about meta-analytic samples, Duval and Tweedie (2000) suggested using a funnel plot to examine publication bias in meta-analysis. In this application, each study's effect size is plotted (e.g., correlation or regression coefficient) along the x-axis against some measure of each study's precision (e.g., standard error) along the

[7] Researchers can refer to the open source codes on GitHub: https://github.com/rdmpage/wikihistoryflow.

FIGURE 4
Multidimensional Scaling Plot Generated Using Country Cultural Distance Data[a]

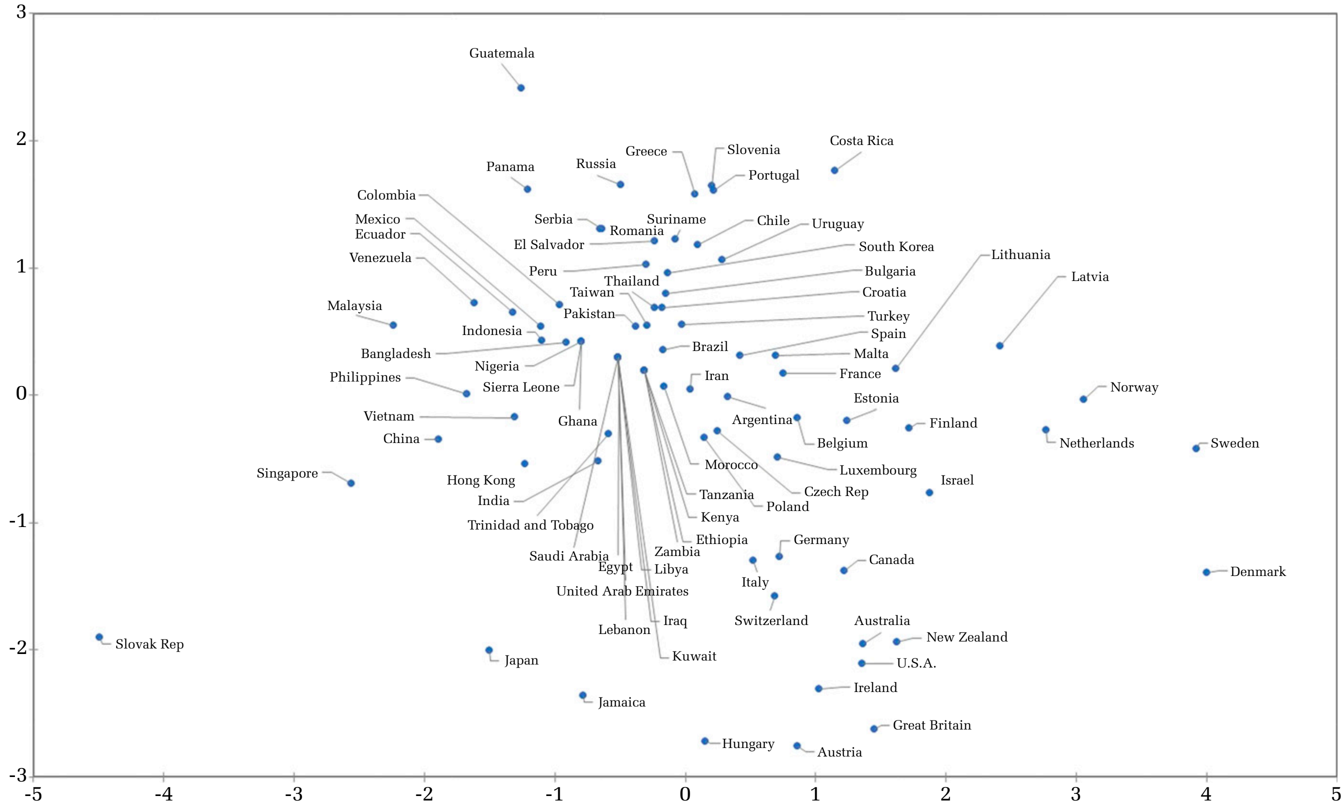

[a] Kogut and Singh (1988) index.

y-axis. In cases of publication bias, a funnel plot will have a skewed and asymmetrical shape.[8]

Recent meta-analytic studies have started to report funnel plots. For example, in their meta-analysis of the consequences of absorptive capacity, Zou, Ertug, and George (2018: 113–115, Appendix 2), reported 10 funnel plots indicating that, while seven of the relationships they meta-analyzed may be influenced by publication bias, three of the relationships were not. Another recent meta-analysis by Joshi, Son, and Roh (2015: 1530, Figure 1) also reports funnel plots to show that there are no strong publication bias issues. Funnel plots can also be used in literature reviews more generally to examine publication bias. For example, Vandenbroucke (1988) generated a funnel plot to show that, while studies that examine the effect of passive smoking on lung cancer do not have a publication bias when the sample is made up of women, those that rely on a sample of men do suffer from such bias.

Caveats. As with many techniques for surveying raw data (e.g., scatter plots), this is a useful tool for recognizing symptoms, but it is not a solution. In addition, not observing manifest indications of publication bias does not guarantee that such bias is absent.

A package called *metafor* in R can be used to conduct a publication check, generate the funnel plot, and also correct for publication bias (Viechtbauer, 2010). If the funnel plot looks asymmetric, providing an indication of publication bias, users can choose to impute new samples to correct for publication bias, so that the resulting funnel plot becomes symmetric.

Maps

A map can be used to visualize the interconnections among countries or regions (e.g., firms in

---

[8] Data visualization is widely employed to examine sample bias by observing the sample distribution. For example, a density plot (a variation of a histogram that uses kernel smoothing to plot values) can visualize the distribution of data. In most cases, the desired shape of the density plot is a normal distribution curve. A skewed density plot indicates that some action should be taken to transform the data to reduce the skewness, or address this in some other suitable way.

FIGURE 5
Map Plotted Using the World Bank's GDP Per Capita Data in 2017

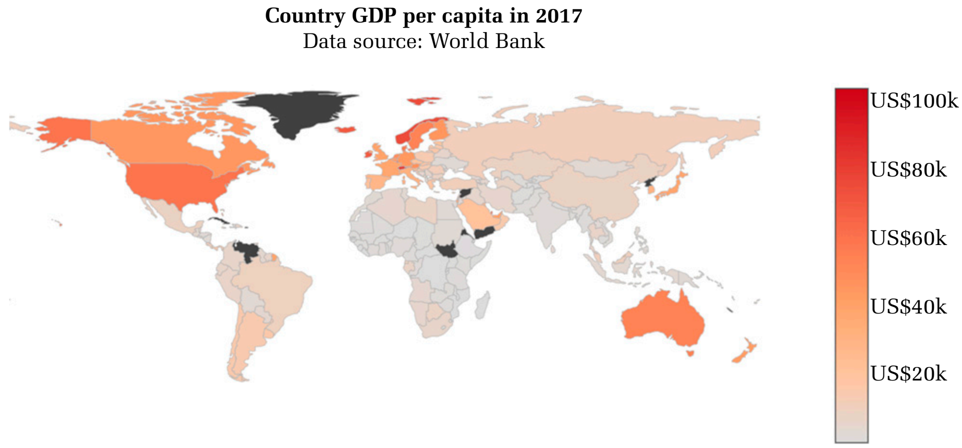

country A entering country B, or patterns of migration) and the differences among countries or regions (e.g., comparison of total volume of foreign direct investments in country A and country B, or investments in research and development in different regions). We can illustrate the interconnections using lines and illustrate the differences using different colors or shades.

Research in international business has employed maps to illustrate differences among geographic areas. For example, Wei and Liu (2006: 552, Figure 1) used a map of China to show that the provinces in mainland China can be broadly divided into three macro areas based on geographic locations. Dheer, Lenartowicz, and Peterson (2015: 447, Figure 1) showed a map of India to indicate that there are nine subcultural regions in India. Research in other areas of science has also used maps. For example, Paez-Espino et al. (2016: 429, Figure 6) used a global map to reveal the global distribution of viral diversity, and Muis, Verlaan, Winsemius, Aerts, and Ward (2016: 3, Figure 1) generated a global map to show the relationship between modeled sea levels and observed sea levels.

Caveats. As with other forms of data visualization, the value of this tool is strongly influenced by the data entered into the tool. Hence, maps can be useful for visualizing special relationships, but visualizing such relationships is not a guarantee of statistically significant relationships.

Plotly's Python graphing package can be used to generate maps.[9] This package supports the world map and many regional maps. Users can choose the color shown on the map and style of the line (e.g., solid or dashed line, color, and width).

In Figure 5, we provide an example of this visualization using country GDP per capita in 2017 (drawing on World Bank data). The darker the red, the higher the GDP per capita in a country. The black areas indicate the countries where the GDP per capita data are not available.

Bubble Plots

Scatter plots are widely used in data analysis to visualize the relationship between two variables. Typically, all the dots in a scatter plot are the same. A "bubble plot" is built on a scatter plot, but is different from most scatter plots because bubble plots have a third dimension, which is the size of each dot. Typically, a bubble plot requires three numerical variables as input: one is represented by the *x*-axis, the other is represented by the *y*-axis, and the third is represented by dot size. For example, when we plot country-level data, the *x*-axis can depict GDP per capita, the *y*-axis can depict average life expectancy, and the size of the dots can represent be the population of the country.

Our brief review suggests that bubble plots have thus far been seldom used in management research, but can be found more frequently in other disciplines. For instance, Nelson, Harvey, Parker, Kastner, Doble, and Gnanapragasam (2013: 7, Figure 3) generated a bubble plot to show cancer detection rates and sample size of each study in three biopsy strategies. The *x*-axis of the bubble plot represented biopsy strategy, the *y*-axis represented cancer detection rates, and the bubble size represented the sample size in each study.

[9] See https://plot.ly/python/maps.

Caveats. Like other data visualization techniques, bubble plots can be an excellent aid for viewers to identify patterns. However, standard disclaimers apply; that is, one has to exercise caution and avoid over-interpreting patterns in randomness.

Bubble plots can be generated using a package called matplotlib in Python. In this package, bubble plots share the same function as a scatter plot, except that users need to configure one argument in the function to represent the size of the bubble. In case bubbles overlap in the plot, users can set the transparency level through the argument of the function. The range of transparency is between 0 and 1, with 0 meaning complete transparency and 1 indicating the most opaque value. Researchers can also use the STATA command *scatter* to produce a bubble plot, where the size of the bubble is specified by the "weight" option of this command.[10]

We provide an example of this visualization in Figure 6. We used open data by Moro, Rita, and Vala (2016) to generate this plot. Each observation in the plot is a post on Facebook. The *x*-axis represents the number of people who clicked on the post, the *y*-axis represents the number of comments on the post, and the bubble size represents the number of people who liked the post.

FIGURE 6
Bubble Plot Generated Using Open Data from Moro et al. (2016)

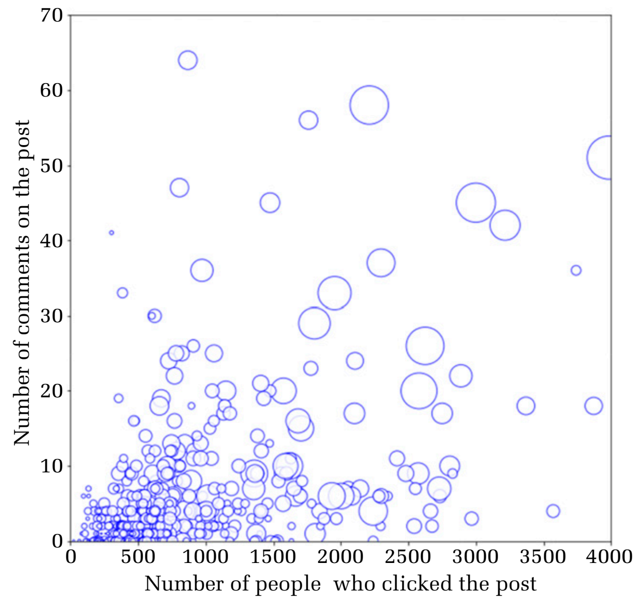

Note: The size of the bubble represents the number of "likes" for the post.

Dynamic Plots

In some cases, researchers might wish to visualize the dynamic processes for a relationship or distribution, to get a sense of how they may change over time. For example, in cases where longitudinal data are available, we may want to consider the pattern in the scatter plot between two variables over time. One way to do this is to generate a scatter plot for each time point, and then compare the different plots. Alternatively, we can generate a dynamic plot by animating these scatter plots, as in a GIF (Graphics Interchange Format) image that you will often see in online conversations.

Research on dynamics is a prevalent topic in management and organization research. For example, research has started to explore the dynamics of social networks (Ahuja, Soda, & Zaheer, 2012).[11] A dynamic plot is particularly useful to visualize the dynamics of social networks (e.g., how actors broker or how the topology of the social network changes).

All the visualization techniques we have discussed in this primer can be animated to illustrate dynamics over time. Additionally, many other basic visualization techniques, such as histograms, bar charts, pie charts, and line charts, can be animated to make them informative with respect to a dynamic process.

Caveats. Visualizing the entire network in a dynamic plot enables researchers to obtain a better understanding of what happens in the network. However, this dynamic visualization might also be distracting when a lot of changes are taking place in the network. Researchers might consider generating separate dynamic plots for networks at the individual level and at the group level, for example.

An open source package called matplotlib in Python helps researchers to produce dynamic plots. It is relatively straightforward to generate a dynamic plot: users first store the various individual plots in a list ("list" is a type of data structure in Python), then use the animation function to show the plots one by one. Users can set the time interval between the plots.

Tree Maps

"Tree maps" display hierarchical data as a set of nested rectangles. Each group is represented by a rectangle, whose area is proportional to its value.

[10] An example of producing bubble plot in STATA can be found here: https://www.stata.com/support/faqs/graphics/gph/graphdocs/scatterplot-with-weighted-markers/

[11] A specific example of a dynamic social network graph is shown here: https://commons.wikimedia.org/wiki/File:Social_graph.gif

FIGURE 7
Tree Map Generated Using Web of Science Citation Data for Bansal and Roth (2000)

Researchers can convert other types of visualization into a tree map. For example, a bar chart can be converted to a tree map, where each rectangle represents a category in *x*-axis of bar chart and the area of the rectangle is proportional to the frequency in *y*-axis of bar chart. Similarly, a word cloud can also be a converted tree map, where each rectangle represents a word and the area of the rectangle is proportional to the frequency of that word.

Tree maps are not commonly used in management and organization research, but have been employed in other research areas. For example, Cable, Ordonez, Chintalapani, and Plaisant (2004: 4, Exhibit 2) generated a tree map to visualize a project portfolio with 41 projects. Each rectangle in the tree map represented a project and the size of the rectangle was proportional to the project budget. Another study by Masood and Lodhi (2015: 54, Figure 2) used a tree map to show the number of variables that cause the ineffectiveness of government audits, where each rectangle in the tree map represented a variable, and, the greater the rectangle size, the more influential that variable.

Caveats. The data for generating tree maps can be single level or multiple level. An example for single-level data is the display of sales figures in different cities; in this case, each rectangle represents a city and the size of rectangle is proportional to sales. An example of two-level data is the display of sales figures in different countries, where the sales figures of different cities are nested in countries; in this case, the big rectangles represent the countries, and the small rectangles represent cities that are nested in the big rectangles, while the size of each rectangle is proportional to the sales of that country or city.

The open source package squarify in Python provides researchers with the ability to generate tree maps. This package can automatically calculate the size of each rectangle and plot it in a way that fits. Users can add a label and color to each rectangle.

We provide an example of a tree map in Figure 7. This tree map depicts the journals that have more than 10 citations to Bansal and Roth (2000). We used citation data from the Web of Science. Each journal is represented by a rectangle and a unique color. The area of each rectangle is proportional to the number of citations to Bansal and Roth (2000) in that journal.

TOOLS FOR DATA VISUALIZATION

Management researchers are often familiar with Excel, SAS, STATA, or SPSS, and it is convenient to use these software packages to produce standard chart types, such as a bar charts, line charts, pie charts, histograms, and scatter plots. Many of these software packages include a user-friendly graphic interface to import data and generate different types of charts.

There are also other software packages that are popular in data science and that allow researchers to generate visually appealing data charts. For instance, Tableau is relatively new but has been growing in popularity. It offers many interactive visualization tools and does a good job with data management. Users can import data from Excel, text files, and database servers. Standard time series charts, bar graphs, pie charts, basic mapping, and so on are supported by this software. R is another programming language and free software environment for statistical computing and graphics. Packages such as ggpot2, plotly, and igraph can be accessed in R to visualize data. As we have noted throughout this primer, Python is also useful and is an interpreted high-level programming language for general-purpose programming. Packages such as matplotlib, seaborn, and networkx can be accessed in Python to visualize data. Compared to Excel, STATA, SPSS, and Tableau, R and Python require users to understand the basics of programming in order to process and visualize their data. That is to say, learning R and Python may take some investment, but they come with the advantage of offering users greater flexibility to customize the visualization of data.

CONCLUSION

Steele and Iliinsky (2010) suggested that beautiful visualizations reflect the qualities of the data they represent, explicitly revealing properties and relationships inherent and implicit in the data that offer insights and new understanding at a glance. The authors further suggested that, for visualization to be beautiful, it should be aesthetically pleasing, novel, informative, and efficient. The purpose of the present editorial was not to provide a detailed step-by-step tutorial of how to perform data visualization. Instead, we hope that it serves as a source of inspiration, an overview, and a guideline to inform readers about some data visualization techniques worth considering. In particular, researchers should reflect not only on what type of visualization should be used, but also how to make such visualization beautiful. Like an artist doing an oil painting, researchers can adjust the size, proportion, color, and luminosity of their visualization to make it more appealing to readers—more pleasurable to look at, easier to understand and to appreciate.

Gokhan Ertug
Singapore Management University

Marc Gruber
Ecole Polytechnique Fédérale de Lausanne (EPFL)

Anthony Nyberg
University of South Carolina

H. Kevin Steensma
University of Washington

REFERENCES

Ahuja, G., Soda, G., & Zaheer, A. 2012. The genesis and dynamics of organizational networks. ***Organization Science***, 23: 434–448.

Arts, S., Cassiman, B., & Gomez, J. C. 2018. Text matching to measure patent similarity. ***Strategic Management Journal***, 39: 62–84.

Bansal, P., & Roth, K. 2000. Why companies go green: A model of ecological responsiveness. ***Academy of Management Journal***, 43: 717–736.

Borgatti, S. P. 2002. NetDraw software for network visualization [Computer software]. Harvard, MA: Analytic Technologies.

Burns, K. 2016. ***Prepared text of the 2016 Stanford Commencement address by Ken Burns*** [Transcript]. Retrieved from https://news.stanford.edu/2016/06/12/prepared-text-2016-stanford-commencement-address-ken-burns

Burt, R. S. 2004. Structural holes and good ideas. ***American Journal of Sociology***, 110: 349–399.

Burt, R. S. 2007. Secondhand brokerage: Evidence on the importance of local structure for managers, bankers, and analysts. ***Academy of Management Journal***, 50: 119–148.

Cable, J. H., Ordonez, J. F., Chintalapani, G., & Plaisant, C. 2004. Project portfolio earned value management using treemaps. Paper presented at PMI Research Conference: Innovations, London, England. Newtown Square, PA: Project Management Institute.

Carney, R. N., & Levin, J. R. 2002. Pictorial illustrations still improve students' learning from text. ***Educational Psychology Review***, 14: 5–26.

Colquitt, J. A. 2013. From the editors: The last three years at *AMJ*—celebrating the big purple tent. ***Academy of Management Journal***, 56: 1511–1515.

Cuéllar, M.-F. 2017. ***Prepared text of the 2017 Stanford Commencement address by Mariano-Florentino***

Cuéllar [Transcript]. Retrieved from https://news.stanford.edu/2017/06/18/prepared-text-2017-stanford-commencement-address-mariano-florentino-cuellar

Cuypers, I. R. P., Ertug, G., Heugens, P. P., Kogut, B., & Zou, T. in press. The making of a construct: Lessons from 30 years of the Kogut and Singh cultural distance index. ***Journal of International Business Studies***.

Dheer, R. J., Lenartowicz, T., & Peterson, M. F. 2015. Mapping India's regional subcultures: Implications for international management. ***Journal of International Business Studies***, 46: 443–467.

Duval, S., & Tweedie, R. 2000. Trim and fill: A simple funnel plot-based method of testing and adjusting for publication bias in meta-analysis. ***Biometrics***, 56: 455–463.

Erkan, G., & Radev, D. R. 2004. Lexrank: Graph-based lexical centrality as salience in text summarization. ***Journal of Artificial Intelligence Research***, 22: 457–479.

Evanschitzky, H., & Goergen, M. 2018. Looking forward, looking back: *British Journal of Management* 2000–2015. ***British Journal of Management***, 29: 3–9.

Fonti, F., & Maoret, M. 2016. The direct and indirect effects of core and peripheral social capital on organizational performance. ***Strategic Management Journal***, 37: 1765–1786.

Gentzkow, M., Shapiro, J. M., & Taddy, M. 2016, July. Measuring polarization in high-dimensional data: Method and application to congressional speech (NBER working paper no. 22423). Cambridge, MA: National Bureau of Economic Research.

Greve, H. 2017. From the editor. ***Administrative Science Quarterly***, 62: v–vi.

Hagberg, A., Swart, P., & Chult, S. D. 2008. ***Exploring network structure, dynamics, and function using NetworkX*** (LA-UR-08-05495; LA-UR-08-5495). Washington, DC: U.S. Department of Energy.

Hofstede, G. 1980. ***Culture's consequences: International differences in work-related values***. Thousand Oaks, CA: SAGE.

Jones, C., Maoret, M., Massa, F. G., & Svejenova, S. 2012. Rebels with a cause: Formation, contestation, and expansion of the de novo category "modern architecture," 1870–1975. ***Organization Science***, 23(6): 1523–1545.

Joshi, A., Son, J., & Roh, H. 2015. When can women close the gap? A meta-analytic test of sex differences in performance and rewards. ***Academy of Management Journal***, 58: 1516–1545.

Kogut, B., & Singh, H. 1988. The effect of national culture on the choice of entry mode. ***Journal of International Business Studies***, 19: 411–432.

Kruskal, J. B. 1964. Multidimensional scaling by optimizing goodness of fit to a nonmetric hypothesis. ***Psychometrika***, 29: 1–27.

Masood, A., & Lodhi, R. N. 2015. Factors affecting the success of government audits: A case study of Pakistan. ***Universal Journal of Management***, 3: 52–62.

McNaught, C., & Lam, P. 2010. Using Wordle as a supplementary research tool. ***Qualitative Report***, 15: 630–643.

Mitra, T., & Gilbert, E. 2014. The language that gets people to give: Phrases that predict success on Kickstarter. In ACM Special Interest Group on Computer–Human Interaction (Ed.), ***Proceedings of the 17th ACM conference on Computer Supported Cooperative Work & Social Computing***: 49–61. New York, NY: ACM.

Moro, S., Rita, P., & Vala, B. 2016. Predicting social media performance metrics and evaluation of the impact on brand building: A data mining approach. ***Journal of Business Research***, 69: 3341–3351.

Muis, S., Verlaan, M., Winsemius, H. C., Aerts, J. C., & Ward, P. J. 2016. A global reanalysis of storm surges and extreme sea levels. ***Nature Communications***, 7: 11969.

Nelson, A. W., Harvey, R. C., Parker, R. A., Kastner, C., Doble, A., & Gnanapragasam, V. J. 2013. Repeat prostate biopsy strategies after initial negative biopsy: meta-regression comparing cancer detection of transperineal, transrectal saturation and MRI guided biopsy. ***PLoS One***, 8: e57480.

Paez-Espino, D., Eloe-Fadrosh, E. A., Pavlopoulos, G. A., Thomas, A. D., Huntemann, M., Mikhailova, N., Rubin, E., Ivanova, N. N., & Kyrpides, N. C. 2016. Uncovering Earth's virome. ***Nature***, 536: 425–430.

Pentland, B. T., Hærem, T., & Hillison, D. 2011. The (n) ever-changing world: Stability and change in organizational routines. ***Organization Science***, 22: 1369–1383.

Robinson, S. L., & Bennett, R. J. 1995. A typology of deviant workplace behaviors: A multidimensional scaling study. ***Academy of Management Journal***, 38: 555–572.

Rothstein, H., Sutton, A. J., & Borenstein, M. 2006. ***Publication bias in meta-analysis: Prevention, assessment and adjustments***. Hoboken, NJ: John Wiley & Sons.

Schmiedel, T., Müller, O., & vom Brocke, J. 2018. Topic modeling as a strategy of inquiry in organizational research: A tutorial with an application example on organizational culture. ***Organizational Research Methods***. Published online ahead of print. doi: 10.1177/1094428118773858

Shaw, J. D., & Ertug, G. 2017. The suitability of simulations and meta-analyses for submissions to *Academy of*

Management Journal. ***Academy of Management Journal***, 60: 2045–2049.

Shepard, R. N. 1967. Recognition memory for words, sentences, and pictures. ***Journal of Memory and Language***, 6: 156–163.

Steele, J., & Iliinsky, N. 2010. ***Beautiful visualization: Looking at data through the eyes of experts***. Sebastopol, CA: O'Reilly.

Vandenbroucke, J. 1988. Passive smoking and lung cancer: A publication bias? ***British Medical Journal (Clinical Research Ed.)***, 296: 391–392.

Viechtbauer, W. 2010. Conducting meta-analyses in R with the *metafor* package. ***Journal of Statistical Software***, 36: 1–48.

Viégas, F. B., Wattenberg, M., & Dave, K. 2004. Studying cooperation and conflict between authors with history flow visualizations. In ACM Special Interest Group on Computer–Human Interaction (Ed.), ***Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*:** 575–582. New York, NY: ACM.

Wattenberg, M., & Viégas, F. B. 2008. The word tree, an interactive visual concordance. ***IEEE Transactions on Visualization and Computer Graphics***, 14: 1221–1228.

Wattenberg, M., & Viégas, F. 2010. Beautiful history: Visualizing Wikipedia. In J. Steele & N. Iliinsky (Eds.), ***Beautiful visualization: Looking at data through the eyes of experts***: 175–191. Sebastopol, CA: O'Reilly.

Wei, Y., & Liu, X. 2006. Productivity spillovers from R&D, exports and FDI in China's manufacturing sector. ***Journal of International Business Studies***, 37: 544–557.

Zou, T., Ertug, G., & George, G. 2018. The capacity to innovate: A meta-analysis of absorptive capacity. ***Innovation***, 20: 87–121.