

PRACTICE FORUM

Augmented intelligence: The new world of surveys at work

Justin Black, Goutham Kurra, and Eric Knudsen*

Glint Inc.–People Science, Redwood City, California, United States

*Corresponding author. Email: eknudsen@linkedin.com

(Received 4 February 2020; revised 14 September 2020; accepted 30 September 2020)

The goal of this article is to inform science, practice, and debate that are related to modern applications of organizational survey research, particularly with respect to the technology-assisted pursuit of happiness and success at work. Our hope is to inspire researchers, practitioners, and HR leaders to start building now for the world of organizational surveys we will live in 15 years from now. The observations and conclusions that are made in this article are drawn from three primary sources: the literature on organizational surveys and employee attitudes, the literature on computational methods, and our personal experience with designing and deploying technology-assisted survey insights at hundreds of large organizations over the past 6 years.

Producing action from organizational surveys remains one of the biggest challenges with them, even though action is generally what drives positive change (Church & Oliver, 2006; Donovan & Brooks, 2006). Frankly, we experts have underperformed on our goal to use surveys to focus and accelerate change, largely due to systems that rely too heavily on administrative involvement, especially from often underresourced human resources (HR) departments/people teams. We believe experts play a critical role in the process (Church & Oliver, 2006) and that expertise can be delivered at greater scale than it is today in terms of both the number of people affected and the size of the effects. Data science, learning and development, and design science are converging to scale personalized insights for leaders, using expert-built algorithms to suggest focus areas and serving up expert-vetted tools and resources through intuitive online experiences.

Given that people will increasingly coexist with machines at work, it is important that our field engages more actively with what we call *augmented intelligence*, or highly specific machine-assisted insights that are faster and smarter than a human could produce. Augmented intelligence is differentiated from artificial intelligence (AI) in that the latter is often operationalized as a broad and complex set of machine-based operations that are organized via an executive function so as to mimic human intelligence. Many models of AI suggest that its purpose is to wholly supplant human intelligence. Even more nuanced models like Searle's (1980) strong and weak AI propose that both the "stronger" and "weaker" forms of AI serve to replace rather than support or augment human intellect. Simply put, augmented intelligence serves to support the human ability to produce valuable insights. In doing so, augmented intelligence augments one small aspect of human intelligence, like the cognitive processing of topics and sentiment from open-ended survey comments, to enrich (rather than fully write) the stories we can tell through data. This article provides the context for and implications of that engagement.

The context for augmented intelligence

People's expectations at work have changed. Employees are impatient with corporate solutions that pale in comparison to the personalized, instant support they get through smartphone apps. There is a notable increase in the popularity of self-awareness and consumerized feedback habits ("thumbs up!"), which makes it natural for employees to expect to influence their work experience or want to know where they stand with respect to their performance and development (Ng *et al.*, 2010). One result is an urgency from both employees and employers to improve performance and build skills. Leaders who wish to proactively steer their teams toward their objectives require faster, more timely, and more relevant feedback and information than what has worked in the past.

The combination of more frequent change and a heightened demand for timely information makes it very difficult to keep a company aligned, and corporations are dissolving faster than ever before (Anthony *et al.*, 2018) despite a movement to embrace digital transformation (Fitzgerald *et al.*, 2013). Data can become stale and irrelevant quickly, including that regarding attitudes (e.g., Petty & Cacioppo, 2012; Wood, 2000). Indeed, across our customers we regularly observe survey scores and comment sentiment changing significantly and sometimes drastically in as little as 3 months.

One benefit of increased focus on digitization is that volumes of data are available in formats that can be readily aggregated. Combined with advances in computing power, digitization is creating opportunities to explore complex within-person relationships across time and experiences. This type of person-centric analysis is firmly grounded in our field's literature (e.g., the three-component conceptualization of organizational commitment introduced by Meyer & Allen, 1991) and is represented in statistical analyses we use in our field like latent profile analysis (e.g., Wang & Hanges, 2011). Unfortunately, until recently, person-centric analyses have been cumbersome and slow to produce useful conclusions. There are many untapped opportunities to leverage data and analytics to generate team- and company-specific insights (Schneider, 2020), but recent developments in alignment between technology and organizational processes have begun to enable unlock such opportunities.

Short- and long-term implications for surveying and reporting

A first step toward augmented intelligence is to adapt people data to meet the needs of a real-time organization to show alignment between leadership practices, people success, and key outcomes. It requires that surveying and reporting be available on demand and that data sources be integrated. With current and emerging survey capabilities, we can begin to explore how the unity of these technologies and methodologies together enrich our ability to generate profound, real-time insights that inspire fundamental improvement in organizations.

On-demand feedback

Well-designed, on-demand feedback programs give employees the opportunity to influence their work experience, their growth, and the business when leaders respond to the feedback in a timely and focused manner. The potential risk of on-demand feedback is that the volume becomes overwhelming. There are at least two common ways to deliver comprehensive and strategic insights (Macey & Schneider, 2006; Schiemann & Morgan, 2006) without overwhelming leaders. One way is to reduce the number of questions asked (e.g., strategic engagement surveys) while increasing the use of open-ended comments, which provide more context not just more volume. Another is to use event-based feedback (e.g., manager effectiveness surveys deployed when a manager reaches 6 months in role, new hire surveys at 30 days), or passive data (e.g., participation in a talent program, time spent engaging with learning modules) to capture specific information in the moment

that can be linked together with other sources at a later time. We have seen both of those tactics work well in practice.

Shortened measures can be just as effective as longer ones. This is especially true in employee surveys like those that are used to measure employee engagement, where many sources of common method variance exist (Podsakoff et al., 2003) and for which the construct of interest is complex and susceptible to that variance becoming bias (Spector, 2006). Prevailing best practice on mitigating bias is to use an a priori approach in which the potential noise is eliminated before the measure is deployed rather than doing so after the results come back (Conway & Lance, 2010). In our own studies, we have found one- and two-item measures of employee engagement to capture up to 90% and 95%, respectively, of the variance in a typical five-item engagement measure. Furthermore, these one- and two-item measures predicted business outcomes with large and statistically significant effect sizes (Glint Inc., 2017). Although these Glint studies investigated the feasibility of short measures for employee engagement, the methods should be investigated for other attitudinal constructs as well. Interestingly, the validity of shortened measures for attitude measurement has also been validated outside of the psychology discipline (e.g., Ang & Eisend, 2018).

Advances in natural language processing (NLP) make it possible to analyze volumes of open-ended comments at a level of accuracy similar to that of human analysts (cf., Speer, 2018). Leaders can instantly see what the team meant when they said they were underrecognized, and how they felt about it (positive, mixed, neutral, or negative), and isolate prescriptive comments to streamline the solution-generation process. This capability—at this level of accuracy and utility—is only recently being introduced in the organizational survey market after decades of research and applications in fields like computational methods and linguistics.

The integration of multiple data sources into a holistic view is not new, but the ease with which this can be done has increased substantially over the past few years. This has moved well beyond multisourced dashboards into cross-source analysis. For example, it is possible to see in a matter of seconds which aspects of the onboarding experience might be correlated with engagement levels at Year 1 for a particular cohort (e.g., female engineers hired for remote positions). The back-end ability to combine, aggregate, and filter data on the fly means that feedback can be sought on the front end for specific reasons and at the time it is most relevant.

On-demand reporting

Well-designed on-demand reporting makes timely information available in personalized ways that all stakeholders and even employees should see them. Timeliness is solved by making results available instantaneously. Although the benefit of rapid analysis is not felt as strongly at the overall level of reporting (i.e., one-second delivery versus one-day delivery is not likely to make or break an action-taking strategy), managers who are empowered to dynamically cut and reanalyze their data across various team views express the greatest influence on rapid analysis and delivery. Doing so eliminates one of the common excuses that leaders give for not taking action on employee engagement data, that the results are “stale” or “too generic.” This type of on-demand reporting is widely available today for organizational surveys.

Personalization is where “reporting” begins to be complemented and ultimately replaced by “augmented intelligence.” Google Maps does not serve up a list of traffic reports, it just tells you which way to go. Some basic versions of this are available now based on formulas that rank survey items based on a combination of descriptive (e.g., difference to benchmark) and correlational (e.g., effect on the employee engagement outcome score) effect sizes, but these solutions are still generic relative to what is possible with augmented intelligence.

On-demand analytics and augmented intelligence

Business intelligence tools and some advanced employee engagement platforms already make it possible to do basic on-demand analytics with people data. More advanced analytics still require significant manual work, particularly when the relationship of interest requires many contextual factors with yet-to-be-measured effects. Thankfully the field is evolving, evidenced by the availability of real-time analytics highlighting specific issues to address or groups¹ of people to support. Although Weiner & McMahon, (2020) covered the use of AI in survey action taking, our focus here is the foundations of machine-assisted insights. Our strong belief is that insights from augmented intelligence are differentiated from past reporting by relevance, timeliness, and scalability.

Insights need to be relevant and timely to be useful. In an age that is characterized by myriad drains on attention, having the right nuggets personalized and adapted to the user is immensely valuable. The kinds of insights that are valuable and actionable by a Human Resources Business Partner (HRBP) are different from those of the CEO. When insights are relevant and timely, the chance of a leader dismissing results as “stale” or “unrelated to my goals” is reduced. But relevance also mitigates the risk of analysis paralysis by surfacing insights in a way that motivates action instead of spurring endless exploration. Consider automatically surfacing to a leader that, given the team’s operational goals this quarter and their ratings and comments about “resources,” they should focus on establishing foundational systems/tools and set fewer goals because they have a tendency to overcommit. These benefits of fast and tailored access are increasingly felt as organizations move to strategies that involve rapid dissemination of results. For example, after deciding to release results to managers right after survey close, LinkedIn saw a 36-point year-over-year increase in the percentage of managers who accessed their results on the first day that they were available. Among senior leaders, they saw a 29-point increase year over year in rapid access. These findings tell a compelling story about the power of easy-to-use technology and instantly available insights.

Augmented insights have a predictive quality that is part machine and part human expert. The utility of a model is determined by how well it can predict the future. Predictive quality, measured by how accurately and quickly the insight can be generated, is achieved by automating a model that is both tuned and then maintained as needed by human experts. It is far more helpful to inform leaders in advance than in retrospect about an attrition issue that stems from a perceived lack of growth opportunities among high-performing women who work remotely.

Although much of the above can be done manually, it is hard to do this at scale. The two types of scale that matter the most to augmented intelligence are the amount of data that are used and the number of people who are using it. Scale increases as the amount of information we collect grows. Survey practitioners will be increasingly working with datasets that contain tens or hundreds of thousands of comments, tied to hundreds of demographic and organizational attributes, across hundreds of point-in-time measurements from surveys, performance feedback, and other passive listening entry points. At this scale, generating relevant and predictive insights on demand to a few administrators is possible with minimal computing power. However, doing so across a large user base so that all stakeholders can have relevant and timely insights personalized to their needs requires smart automation (e.g., selective processing) and human-taught machine intelligence that adapts to different user types (executives, managers, individual contributors, etc.).

Data architecture for augmented intelligence

To deliver on the need for insights from augmented intelligence to be relevant and timely, quickly and accurately predictive, and scalable as the amount of data and users grow, we require fundamentally different ways to represent data. It starts with the data models.

¹In this context a group could be a node in the organizational hierarchy, a project team, or any set of individuals who share common characteristics as defined by attributes in the HR record (e.g., new hires who work remotely and part time).

We call today's standard data model, a legacy from the past, a "survey-centric" architecture. In this architecture, the survey is the center of the data universe. Surveys have questions, and responses are collected for each question. The responses are tied to individuals whose demographics and other attributes are then used to slice and dice the response data to produce insights while maintaining a level of confidentiality. The drawback of this architecture is that surveys end up in silos. It becomes hard to tie the data together so that insights can be generated across different surveys or channels of feedback, cut by other data points.

In contrast, a "people-centric" architecture is a data representation (sometimes called a "graph database") that allows for individuals to be the center of the universe. Survey responses, transactional human-resource-information-system (HRIS) events (e.g., a change in role), signals from other channels (e.g., applicant tracking systems), demographics, and other metadata are captured on an individualized, versioned timeline. The metadata may include references to the source of the data, mechanisms of collection, confidentiality, relevant timestamps, survey question order, and survey constraints (e.g., min and max choices for scalar items). Multiple graph structures, such as organizational hierarchies and relationships between individuals, can be superimposed on the individuals or, for even more flexibility, over individual data points. A well-built people-centric technology will not put the onus on individual users (e.g., HR admin) to join all of this data themselves but will seamlessly accept people data from other systems or generate its own metadata.

People-centric representations allow for quick aggregations of people along multiple axes while properly maintaining confidentiality. For example, say an organization runs a nonconfidential exit survey and a semiannual engagement survey that is confidential with reporting allowed to groups of five or more respondents. If we wanted insight into why people who were previously engaged at work are leaving the company, this would be very difficult to do with a survey-centric architecture. Such reporting typically requires special data warehouses or error-prone manual data management. Not only is this time consuming; confidentiality can be compromised as well. On the other hand, a people-centric data model makes it easy to not only aggregate all people who are favorable on engagement dimensions in a survey 6 months prior to their exit interview and extract their departure reasons but also maintain confidentiality. Note that because the confidentiality of each piece of data that is being captured resides as metainformation with the data itself, such combinations of information can apply the strictest of multiple thresholds or a function designed to keep the degrees of freedom of the information the same or higher.

There is significant context that can be captured either as data or metadata along people-centric timelines that allow for rich sets of insights and actions to result. For instance, longitudinal studies that involve cohorts are made very easy, especially in real time. Interventions and control group experiments can be conducted and analyzed in a fraction of the time. Complex multivariate cross-lagged analyses can be largely automated using machine learning in order to predict behavioral and business outcomes—and then identify where to intervene and on what issues. Key considerations like confounds and levels of analysis (Harter & Schmidt, 2006) can be automatically applied to analyses. Organizational network metadata showing how and with whom we communicate, a commonly available source of passive data, can be incorporated into everything from automated feedback-seeking to personalized action-taking recommendations. Expert-created typologies can be instantaneously compared and nested with unstructured text to tell us what people meant when they said they were unhappy with their career opportunities.

Computational methods and data models

Advances in computational methods and statistical modeling have allowed for very efficient methods for doing complex analyses in real time (Vicknair et al., 2010) and also for doing large-scale automated machine learning. One key element of these advances is an approach called "parallelization," whereby a large task is virtually delegated across multiple processors. Imagine multiple people working to calculate one mean of millions of data points: Each person takes a chunk of data

and performs the same step on their chunk before bringing them back together to get the final average. Rather than multiple people, multiple processors (sometimes in the same computer, sometimes in different ones) can be used to automatically operate “in parallel” to support processing large datasets extremely quickly.

In a people-centric model, each unit is a person who is associated with demographics, timelines of events, feedback, responses, or even other individuals (generating a graph structure). For example, one dataset might have 50,000 employees of a manufacturing organization that includes ethnicity, department name, tenure, a behavioral outcome like safety incident frequency, and survey data across multiple programs (onboarding, engagement, and safety culture). Along with each employee, there may be HRIS transactions spanning multiple years (e.g., promotions, training, and changes in reporting relationships). Finally, one might have both numerical responses and comment data (e.g., topics, sentiment, prescriptive phrasing, and favorability all preextracted from the text). The end result is a massive dataset that is built around the individual, and parallelization provides a fast framework for using highly sophisticated algorithms with these data on shorter timelines. Better yet, with the advent of cloud-computing services, many of these complex computing processes can be outsourced to platforms that handle parallelization without human intervention. Finally, survey vendors can build products directly on top of these cloud technologies and remove the need for client organizations to build such “tech stacks” in house.

Regarding algorithms, most modern machine-learning approaches fall into two categories: supervised and unsupervised methods. Supervised methods are typically predictive in nature, meaning they are used to produce a final value based on preexisting data. This final value might be a label (e.g., should the employee termination reason “Unhappy with my pay” be labeled as related to “Compensation” or “Team dissatisfaction?”) or a precise numeric value (e.g., the probability of an employee leaving is 65%). In a process similar to that prescribed in our field for validating predictors (cf. Cohen *et al.*, 2003), these supervised methods produce a validated model that is built using input data (say, two historical years of safety data) and is tested for predictive accuracy on an “unseen” pool of data (e.g., a third year of safety data; Gholamy *et al.*, 2018). On the other hand, unsupervised methods are often used to impose structure on data rather than produce a predicted value or label. One common unsupervised method is “clustering,” which is an approach for surfacing natural groups that emerge in the data (e.g., Are there unique, unseen patterns of employee responses to an engagement survey that might warrant action planning differently for these “groups”?).

The combination of these statistical methodologies with advances in technology like parallelization means that massive data sets can be scanned for statistical anomalies or relationships and those relationships can be further tested against held-out data. This process is like having a team of data analysts working to find deeply interesting insights and only surfacing the most relevant and accurate insights, automatically. Interestingly, these algorithms can be further optimized through a process called “tuning” to boost insights that are linked to specified outcomes of interest (e.g., turnover, safety), generating highly personalized insights. Furthermore, multiple insights can be statistically linked together in ways that drive clearly to specific action. For example, a linked insight that could conceivably be generated automatically is “Did you know that factory workers in the following jobs have 60% fewer injuries when they report being satisfied with their training and that training satisfaction is highly correlated with the completion of course X for workers with tenure < 3 years, but course Y for workers with tenure \geq 3 years?”

Data protection and privacy

In order to be compliant and ethical across the world, we must deliver augmented intelligence and its personalized insights in a way that protects the confidentiality of people giving feedback, and a people-centric architecture enables that. We believe the field can go beyond compliance, not only to protect users but also to proactively enhance their experience through language- and

culture-based personalization that makes people feel valued. Because the data are centered around the individual and metadata can be used to manage confidentiality and access permissions, a well-designed data architecture democratizes access to insights while preserving commitments to respondents' data and privacy. With a people-centric architecture, it could even be possible for companies to facilitate the portability of peoples' own data so that they can bring their data with them to new organizations without violating company or individual data protections. Some large technology companies like Facebook have already begun to offer such functionality in the wake of increased public scrutiny regarding personal data.

Augmented intelligence and employee engagement

The employee engagement industry is positioned to benefit greatly from augmented intelligence, as more companies are increasing the frequency of engagement feedback. The shift to frequent and more integrated feedback was driven by demand from business leaders—not by any revolutionary scientific findings. This signals an important shift in our field, from expert-determined to user-demanded solutions, that nonetheless must be rooted in good science and practice. Managers have been educating themselves on how to be great people leaders more so than in the past (Riggio, 2008); the tools we give them need to keep pace with their skills.

Like many other innovations, novel approaches to employee engagement have been built out of necessity and around specific constraints. The first constraint was time. Chief human resource officers have a difficult job justifying investments in reporting tools that do not provide data in real time or on demand, which has become a basic standard in business reporting. In order to compete, many engagement solutions began serving up results instantaneously, increasing speed of access. People-centric architectures and augmented intelligence create the velocity that is necessary to enable the next phase: speed of insight. Here are a few examples of efficiency and specificity gains possible with current technology:

- Dynamic HRIS updating, making it possible to change the present view of results to match a historical view of the organization and to change the historical view of results for the purposes of trend. For example, one could toggle between a view of the trend based on “following the manager” and one based on “following the employees.”
- Multimatrixed hierarchy reporting superimposed on multiple hierarchies, which enables leaders in a matrixed organization to see the results for all business units, functional groups, or teams for which they have responsibility.
- Automatic identification of groups with an increased/decreased probability of doing something we care about (e.g., selling more, staying longer) because of something we can control (e.g., fit in role, goal clarity). We do not advocate the use of this technology to identify *individuals* unless it is part of an opt-in program (Saari & Scherbaum, 2020). The risks associated with identifying individuals are too great to justify the benefit. Interventions for engagement are more than adequately executed at the group level.
- Automated analysis of open-ended comments, which can be analyzed instantaneously, enabling a leader in minutes (instead of hours) of reading to understand the context and nuances of a particular issue and isolate employees' recommendations for improvement.

The second constraint was frequency. Once the data-to-insight-to-action lag was shortened through real-time reporting, users demanded more frequent insights about their people, to align with their other business reporting and to assess progress throughout the year. This constraint was more challenging because it required a departure from the traditional 50–75 item engagement survey. On the survey design side, companies responded in a few different ways:

1. Survey everyone frequently by reducing survey lengths by using statistical item reduction and strategic selectivity, and take advantage of comment data analyzed instantaneously.
2. Survey equal parts of the population periodically (e.g., 25% per quarter).
3. Survey everyone periodically and run sample surveys between these census surveys.

We prefer the first approach because it enables all teams to get a full update on progress each survey and it avoids common pitfalls of sampling (Mastrangelo, 2020). The other revelation driven by the frequency constraint was that action taking had to be simplified (Weiner & McMahon, 2020); “big event” approaches to feedback utilization were too much work to do regularly. One of the most efficient ways to encourage the use of feedback is to work it into an existing workflow or structure for regular conversations at work. Augmented intelligence is well positioned to do that today, and increasingly so over the next few years. Already today, engagement solutions use algorithms to show leaders strengths, opportunities, and suggested actions. This next evolution of augmented intelligence will involve personalized guidance, further putting “the coach in the machine,” with the goal of improving the quality of conversations people have at work. In doing so, augmented intelligence is and will continue to shift (not replace) an HR business partner’s responsibility from “analyst” to “advisor,” reserving valuable time and energy for the most complex, most human of situations.

As conversations become the center of people success programs (versus surveys or assessments), the silos between attraction, engagement, performance, and learning will go away. Instead, engagement survey results will be just one of a set of factors that inform regular conversations through personalized insights, recommendations, and motivational cues, all of which are instantaneously integrated, analyzed, and delivered by augmented intelligence.

Conclusion

It’s time to embrace augmented technology because it offers an opportunity to increase the scale the influence of our field and bring humanity to the world of work. We have the ability to use it effectively, but we must also use it responsibly. This requires both technology and subject matter expertise across engineering, data and people scientists, and end users. Our success as a field will be greater if we can let go of old approaches for which better ones now exist. Agile feedback methodologies such as frequent engagement pulsing or life-cycle programs have not been extensively reported on in the academic literature. We hope the demand for augmented intelligence will inspire deeper understandings of how to optimize the timeliness, relevance, predictive quality, and scalability of survey insights. Our field is perfectly suited to pioneer that research.

References

- Ang, L., & Eisend, M. (2018). Single versus multiple measurement of attitudes: A meta-analysis of advertising studies validates the single-item measure approach. *Journal of Advertising Research*, 58(2), 218–227.
- Anthony, S. D., Viguier, S. P., Schwartz, E. I., & Van Landeghem, J. (2018). 2018 corporate longevity forecast: Creative destruction is accelerating. Innosight. <https://www.innosight.com/wp-content/uploads/2017/11/Innosight-Corporate-Longevity-2018.pdf>
- Church, A. H., & Oliver, D. H. (2006). The importance of taking action, not just sharing survey feedback. In A. I. Kraut (Ed.), *Getting action from organizational surveys: New concepts, technologies, and applications* (pp. 102–130). Jossey-Bass.
- Cohen, J., Cohen, P., West, S. G., & Aiken, L. S. (2003). *Applied multiple regression/correlation analysis for the behavioral sciences* (3rd ed.). Lawrence Erlbaum Associates.
- Conway, J. M., & Lance, C. E. (2010). What reviewers should expect from authors regarding common method bias in organizational research. *Journal of Business and Psychology*, 25, 325–334.
- Donovan, M. A., & Brooks, S. M. (2006). Leveraging employee surveys to retain key employees: A means to an end. In A. I. Kraut (Ed.), *Getting action from organizational surveys: New concepts, technologies, and applications* (pp. 456–482). Jossey-Bass.

- Fitzgerald, M., Kruschwitz, N., Bonnet, D., & Welch, M.** (2013). Embracing digital technology: A new strategic imperative. *MIT Sloan Management Review*, *55*, 1–12.
- Gholamy, A., Kreinovich, V., & Kosheleva, O.** (2018). *Why 70/30 or 80/20 relation between training and testing sets: A pedagogical explanation.* (TEP-CS-18-09). University of Texas at El Paso. https://scholarworks.utep.edu/cs_techrep/1209.
- Glint Inc.** (2017, June 7). New data reveals direct link between employee engagement and Glassdoor scores, stock market performance [Press release]. <https://www.glintinc.com/press/new-data-reveals-direct-link-employee-engagement-glassdoor-scores-organizational-performance/>
- Harter, J. K., & Schmidt, F. L.** (2006). Connecting employee satisfaction to business unit performance. In A. I. Kraut (Ed.), *Getting action from organizational surveys: New concepts, technologies, and applications* (pp. 33–52). Jossey-Bass.
- Macey, W. H., & Schneider, B.** (2006). Employee experiences and customer satisfaction: Toward a framework for survey design with a focus on service climate. In A. I. Kraut (Ed.), *Getting action from organizational surveys: New concepts, technologies, and applications* (pp. 53–75). Jossey-Bass.
- Mastrangelo, P. M.** (2020). Improving the design and interpretation of sample surveys in the workplace. In W. H. Macey & A. A. Fink (Eds.), *Employee surveys and sensing: Driving organizational culture and performance* (pp. 38–52). Oxford University Press.
- Meyer, J. P., & Allen, N. J.** (1991). A three-component conceptualization of organizational commitment. *Human Resource Management Review*, *1*, 61–89.
- Ng, E. S., Schweitzer, L., & Lyons, S. T.** (2010). New generation, great expectations: A field study of the millennial generation. *Journal of Business and Psychology*, *25*, 281–292.
- Petty, R., & Cacioppo, J. T.** (2012). *Communication and persuasion: Central and peripheral routes to attitude change.* Springer Science & Business Media.
- Podsakoff, P. M., MacKenzie, S. B., Lee, J.-Y., & Podsakoff, N. P.** (2003). Common method biases in behavioral research: A critical review of the literature and recommended remedies. *Journal of Applied Psychology*, *88*, 879–903.
- Riggio, R. E.** (2008). Leadership development: The current state and future expectations. *Consulting Psychology Journal: Practice and Research*, *60*, 383–392.
- Saari, L., & Scherbaum, C.** (2020). From identified surveys to new technologies: Employee privacy and ethical considerations. In W. H. Macey & A. A. Fink (Eds.), *Employee surveys and sensing: Driving organizational culture and performance* (pp. 391–406). Oxford University Press.
- Schiemann, W. A., & Morgan, B. S.** (2006). Strategic surveys: Linking people to business strategy. In A. I. Kraut (Ed.), *Getting action from organizational surveys: New concepts, technologies, and applications* (pp. 76–101). Jossey-Bass.
- Schneider, B.** (2020). Strategic climate research: How what we know should influence what we do. In W. H. Macey & A. A. Fink (Eds.), *Employee Surveys and Sensing: Driving organizational culture and performance* (pp. 121–134). Oxford University Press.
- Searle, J. R.** (1980). Minds, brains, and programs. *Behavioral and Brain Sciences*, *3*(3), 417–424.
- Spector, P. E.** (2006). Method variance in organizational research: Truth or urban legend? *Organizational Research Methods*, *9*, 221–232.
- Speer, A. B.** (2018). Quantifying with words: An investigation of the validity of narrative-derived performance scores. *Personnel Psychology*, *71*, 299–333.
- Vicknair, C., Macias, M., Zhao, Z., Nan, X., Chen, Y., & Wilkins, D.** (2010, April 15–17). *A comparison of a graph database and a relational database: A data provenance perspective* [Paper presentation]. Proceedings of the 48th annual Southeast regional conference, Oxford, MS, United States. 10.1145/1900008.1900067
- Wang, M., & Hanges, P. J.** (2011). Latent class procedures: Applications to organizational research. *Organizational Research Methods*, *14*, 24–31.
- Weiner, S. P., & McMahon, M.** (2020). Action taking augmented by artificial intelligence. In W. H. Macey & A. A. Fink (Eds.), *Employee surveys and sensing: Driving organizational culture and performance* (pp. 338–354). Oxford University Press.
- Wood, W.** (2000). Attitude change: Persuasion and social influence. *Annual Review of Psychology*, *51*(1), 539–570.