



Contents lists available at ScienceDirect

Journal of Business Research

How crowdsourcing improves prediction of market-oriented outcomes[☆]Mark Lang^{a,*}, Neeraj Bharadwaj^{b,1}, C. Anthony Di Benedetto^{c,2}^a Saint Joseph's University, 303 Mandeville Hall, 5600 City Avenue, Philadelphia, PA 19131, USA^b University of Tennessee, 323 Stokely Management Center, 916 Volunteer Boulevard, Knoxville, TN 37996-0530, USA^c Fox School of Business, Temple University, 1801 Liacouras Walk, A523 Alter Hall (006-09), Philadelphia, PA 19122-6083, USA

ARTICLE INFO

Article history:

Received 27 December 2014

Received in revised form 7 March 2016

Accepted 8 March 2016

Available online xxx

Keywords:

Crowdsourcing

Market information processing

Market prediction

Forecasting

Organizational learning

ABSTRACT

Firms often struggle to be proficient in predicting uncertain market conditions and forecasting the outcomes of their business initiatives. This research introduces crowdsourcing as an innovative tool that can enhance market information processing, and in turn, improve prediction of market-oriented outcomes (e.g., sales). We field test a forecasting tournament with employees at a Fortune 100 consumer packaged goods firm, and examine the extent to which predictions based on the “wisdom of the crowd” outperform those generated by traditional forecasting approaches. We find that crowdsourcing produces results superior to the firm’s incumbent approaches almost three-fourths of the time across a broad range of business decisions. Additionally, we conduct a survey with participants to open up the “black box” of crowdsourcing. We find that differences in information acquisition and interpretation are the underlying mechanisms that can explain the improved prediction accuracy found through crowdsourcing.

© 2016 Elsevier Inc. All rights reserved.

1. Introduction

Managers must continually respond to a rapidly changing marketplace in order to protect their competitive advantage. This obligation is recognized across diverse disciplines such as marketing (Day, 2011), innovation (Bharadwaj & Noble, 2015), strategic management (D’Aveni, Dagnino, & Smith, 2010), and forecasting (Armstrong, 2006). Dynamic capabilities, a firm’s ability to adapt competences and resources to better respond to the marketplace (Eisenhardt & Martin, 2000; Teece, Pisano, & Shuen, 1997), play an important role in this regard. As business leaders manage this evolution, one of the most important and difficult tasks they undertake is to predict uncertain market conditions and future business outcomes (Day, 2011; Yan & Ghose, 2010). The ability to make effective predictions about markets is especially important as managers develop strategies and implement plans, and forecast the resulting impact on sales, profit, and firm value (Morgan, 2012; Rao & Bharadwaj, 2008). This article introduces crowdsourcing as an innovative decision support tool that can enhance market information processing and directly improve market-oriented

predictions (e.g., consumer preferences or competitor response) and business forecasts (e.g., sales or profits) (Narver & Slater, 1990).

These market-oriented predictions occur at the intersection of important dynamic capabilities: market learning is translated into forecasts that support planning and implementation, new product development, pricing, and strategic decision-making (Morgan, 2012; Vorhies & Morgan, 2005). Market learning not only plays a critical role at this intersection (Cepeda & Vera, 2007) but it is also identified as one of the most significant areas of dynamic and marketing capabilities improvement (Morgan, 2012; Vorhies & Morgan, 2005; Vorhies, Orr, & Bush, 2011).

Unfortunately, forecasting often struggles to be proficient in forecasting (Kahn, 2002; Srivastava, Shervani, & Fahey, 1999; Yan & Ghose, 2010). For example, in response to a major product launch failure, attributed largely to forecasting errors, Procter & Gamble’s CEO pledged to make better use of online tools to support demand forecasting and innovation initiatives (Neff, 2012). As an example of the importance of forecasting in designing and executing business strategies, the Marketing Science Institute designates providing guidance to firms on forecasting as one of its foremost 2014–2016 research priorities (MSI, 2014). These examples raise the question of how firms can improve market-oriented predictions and forecasts.

The foundational research establishes that improvements in the development and use of market learning can be achieved through improvements to the market information processing that produces it: the acquisition, distribution/transmission, interpretation/processing, utilization, and storage of market information in organizational memory (Day, 1994; Moorman, 1995; Sinkula, 1994).

[☆] This research benefitted from the comments and guidance of J. Jeffrey Inman, Nathan M. Fong, Nicolle Clements, and participants attending the presentation at the 2014 American Marketing Association Winter Educators’ Conference.

* Corresponding author. Tel.: +1 610 660 3431.

E-mail addresses: mlang@sju.edu (M. Lang), nbharadwaj@utk.edu (N. Bharadwaj), anthony.dibenedetto@temple.edu (C.A. Di Benedetto).

¹ Tel.: +1 865 974 1610.

² Tel.: +1 215 204 8147.

Better use of decision support tools is identified as one way to achieve improvements in forecasting (Diamantopoulos & Winklhofer, 2003; Krishnan & Ulrich, 2001; Wierenga, Van Bruggen, & Staelin, 1999). Vorhies and Morgan (2005), for example, refer to the potential of market-based learning tools to improve capabilities such as market information management. In this regard, the forecasting literature advocates exploring innovative techniques to improve market-oriented predictions (Armstrong, 2006). One of the emerging techniques in business and in other disciplines is crowdsourcing. Most commonly, crowdsourcing applications use a central technology platform to collect and integrate information and the opinions of a diverse set of individuals to predict uncertain future outcomes, solve problems, or design solutions (Servan-Schreiber, 2013; Surowiecki, 2005). Recent writings suggest that these applications can serve as effective decision support tools directly satisfying the information processing improvement needs described in earlier writings (Armstrong, 2006; Bonabeau, 2009; Malone, Laubacher, & Dellarocas, 2010). We propose that crowdsourcing applications can serve to improve market information processing and market-oriented predictions and forecasts.

Forecasting and decision making are often classified as quantitative-data driven or management knowledge-judgment driven (Wierenga et al., 1999). Crowdsourcing applications, by pooling a wide number of opinions in a structured way, combine qualities of both of these types of forecasting methods described above.

Despite the emerging use of crowdsourcing to forecast political science, financial, economic, and legal outcomes (Arrow et al., 2008), the wisdom of the crowd has yet to be widely adopted in business and marketing. There has been limited empirical testing and, more importantly, virtually no explanation as to how crowdsourcing platforms produce superior results. This fact is surprising given the well-defined need for innovative tools to improve predictions of important market-oriented outcomes (Day, 2011). Before they can be considered as a new methodology, however, their efficacy must be demonstrated and their mechanism explained (Armstrong, 2006).

Given that the earlier studies have not empirically investigated *how* crowdsourcing works, both researchers and practitioners have declared the need to open up the “black box” (Green, Armstrong, & Graefe, 2007). Armstrong (2006), for example, performs a detailed review of traditional and emerging forecasting techniques and calls for further use and testing of crowdsourcing applications such as prediction markets. He suggests testing these new techniques against other traditional methods in the realm of business, and observes that there has been little research into their efficacy. Matzler, Grabher, Huber, and Füller (2013) is an example of work in this direction.

This issue leads to the primary research questions addressed in this paper: can crowdsourcing applications improve the processing and use of market information to produce more accurate predictions, and if so, how? Accordingly, the objectives of this study are to: (i) empirically test the extent to which crowdsourcing can improve predictions relative to the methods traditionally employed by firms and, more importantly, (ii) examine how they produce superior results. To address these questions, we present a non-trading-based, crowdsourcing tournament as performing market information processing in a uniquely adapted way, applied to specific market-oriented business questions.

We design and implement a field test at three separate divisions of a Fortune 100 consumer packaged goods firm, with emphasis on predicting a wide range of market-oriented outcomes such as new product sales, supply chain shipments, and expected dollar sales (Srivastava et al., 1999). Our results reveal that the crowdsourcing application is more accurate than the firm's incumbent forecasting methodologies 73% of the time, reduces average error by almost 34 percentage points, and reduces the error range by over 40%. In examining the crowdsourcing mechanism associated with their performance, we find differences in *information acquisition* (i.e., the extent to which participants search for additional information to use in their individual

predictions) and in *interpretation* (i.e., the extent to which participants bring different perspectives and thinking to their prediction) to be directly and positively related to improvements in prediction accuracy. This research thereby extends the literature by demonstrating the efficacy of crowdsourcing, and providing insight into why the “wisdom of the crowd” can yield more accurate forecasts.

2. Literature review

2.1. Crowdsourcing applications

Crowdsourcing applications aggregate the judgment of many people across the firm to predict uncertain market outcomes, the resulting aggregated prediction may be more accurate than traditional methods used within firms because with crowdsourcing, individual participants process market information differently. We propose that the way acquisition, distribution, and interpretation are performed within a crowdsourcing application produces superior market prediction results. As a result, these applications may perform market information processing better than traditional methods commonly used within firms.

Crowdsourcing applications can be categorized into trading and non-trading-based platforms. The former include, for example, decision markets, political stock markets, event futures, and prediction markets (Horn, Ohneberg, Ivens, & Brem, 2014; Servan-Schreiber, 2013). These have been used to predict political, finance, movie, and sports outcomes. In business, this type of application has been deployed at firms such as Hewlett-Packard (Ho & Chen, 2007), Intel (Hopman, 2007), and Google (Cowgill, Wolfers, & Zitzewitz, 2009) where they have been found to perform as well as, and in many cases, better than existing traditional methods in terms of estimates, error rates, and variability. For example, Hopman (2007) and Ho and Chen (2007) report more accurate predictions in 75% of the cases examined.

For reasons related to trading complexity with trading-based platforms, non-trading-based crowdsourcing platforms—which are easier for people with no stock trading experience to use and do not suffer from legal challenges associated with betting—are emerging (Servan-Schreiber, 2013). These include, for example, forecasting competitions or tournaments (Afuah & Tucci, 2012; Good Judgment Project, 2014) and idea markets (Soukhoroukova, Spann, & Skiera, 2012), which are more suitable for use in corporate settings.

The superior accuracy of crowdsourcing has been attributed to its ability to collect and combine widely dispersed information and opinions from a diverse group of participants, focus them on specific marketing questions, and have them predict a single common outcome. They efficiently and accurately aggregate, weight, and average information from numerous participants (Ho & Chen, 2007; Hopman, 2007). These tasks are accomplished continuously, in real-time through a central, interactive technology platform.

2.2. Crowdsourcing and market information processing

As noted earlier, improvements in market learning can be achieved through improvements to the market information processing that produces it. The literature describes market information processing as comprising the acquisition, distribution, interpretation, utilization, and storage of market information in organizational memory (Day, 1994; Moorman, 1995; Sinkula, 1994). Interestingly, acquisition, distribution, and interpretation are also found across the disparate crowdsourcing and collective intelligence literatures as design elements that are suggested as producing superior prediction accuracy (Bonabeau, 2009; Malone et al., 2010). These factors have not yet been tested empirically in these fields. Based on these findings, we propose that in a crowdsourcing application, key components of market information processing (*acquisition, distribution, and interpretation*) can be

performed in a uniquely adapted way that can be applied directly to specific marketing questions and decisions.

To this end, we examine the impact on accuracy of differences in acquisition, distribution, and interpretation between a crowdsourcing application and existing forecasting methods being used within a subject company. In the ensuing discussion, information acquisition, distribution, and interpretation are defined and explained specifically as *individual differences in acquisition*, *use of shared information*, and *differences in interpretation*. These are presented in Section 2.2.1 as an overlapping set of market information processing and crowdsourcing factors that work to improve prediction accuracy for specific tasks. Measurement scales for these factors are presented in Appendix A.

2.2.1. Acquisition

Information acquisition will refer to the actions that participants in the crowdsourcing tournament take to obtain additional information to use in their individual predictions. This information can be specialized data and observations, new or existing, and external or internal information (Moorman, 1995; Sinkula, 1994; Sinkula, Baker, & Noordewier, 1997). The acquisition of new information by participants facilitates the development of richer mental models regarding a given phenomenon, which in turn, promotes a more nuanced understanding of a decision situation (Day & Nedungadi, 1994). The richer set of associations possessed by the firm's employees can affect the breadth of organizational learning (Huber, 1991; March, 1991), which has been noted to improve forecasting specifically (Kahn, 2002; March, 1991; Slater & Narver, 2000). We advance the following hypothesis:

H1. The more market information participants acquire to make their predictions, the higher will be their prediction accuracy.

2.2.2. Distribution

Distribution of information is also described as an important dimension of information processing (Maltz & Kohli, 1996; Moorman, 1995; Sinkula et al., 1997). The crowdsourcing technology platform distributes two forms of information: the first is background information that establishes a minimum baseline of common information for all participants. The second is the ability for participants to view aggregated predictions made by other participants in real time. Distribution, in this context, is represented as participants accessing this information shared in the prediction competition to make and update their own predictions. Accessing shared information can lead to new and refined insights and understanding and lead to improved decisions and actions (Huber, 1991; Sinkula, 1994; Slater & Narver, 1995). In the crowdsourcing literature, publicly viewable information and results: transfer private information from knowledgeable informants to less knowledgeable informants, allow participants to learn from others, update their own information, and update their predictions (Ho & Chen, 2007; Hopman, 2007). We hypothesize the following:

H2. The more participants use shared information to make their predictions, the higher will be their prediction accuracy.

2.2.3. Interpretation

Before an organization can act on new information, it must be interpreted or given meaning (Sinkula, 1994). We advance here that differences in interpretation equates to participants using any combination of different mental models, cognitive approaches, processes, analyses, opinions, and assumptions to make their prediction (Day, 1994; Moorman, 1995). These differences result from the crowdsourcing application's ability to assemble a group of participants (from across the organization) that possess different experiences, tenure, training, expertise, perspectives, skills, or abilities (Milliken & Martins, 1996; Rodan & Galunic, 2004). These may come, for example, from people in different hierarchical levels, functional areas,

organizational cultures, geographical areas, or network roles (Day & Nedungadi, 1994; Servan-Schreiber, 2012).

Exposing new information to multiple interpretations and challenging assumptions improves results by expanding learning (Sinkula, 1994; Slater & Narver, 1995), introducing non-redundant perspectives (March, 1991), and reducing the negative effects of group interactions such as groupthink or dominant ideology (Day, 2011; Surowiecki, 2005). Specific to forecasting, Armstrong (2001) finds that "the more that data and methods differ, the greater the expected improvement in accuracy over the average of the individual forecasts" (p. 2). We hypothesize the following:

H3. The greater the difference in interpretations that participants use to make their predictions, the higher will be their prediction accuracy.

3. Methodology and study design

A field study was implemented to examine if and how crowdsourcing applications may improve the processing and use of market information in support of better marketing predictions. As part of an academic-practitioner research collaboration, a custom-designed crowdsourcing tournament was implemented within a Fortune 100 company. The outcomes of a series of crowdsourcing tournaments were compared against corresponding forecasts produced from traditional methods within the subject company. After the crowdsourcing tournament, a survey was sent to participants to examine specifically how crowdsourcing may have improved market information processing. Lumenogic, a recognized leader in crowdsourcing, assisted with the design and hosted the crowdsourcing technology platform for this study. The crowdsourcing tournament ran within three autonomous divisions of the U.S. operations of a consumer packaged goods food company operating in three different product categories. It is a publicly traded, mature company, operating in well-developed industry.

3.1. Crowdsourcing tournament

Although many crowdsourcing implementations engage participants from both inside and outside an organization, the host firm's requirements for confidentiality drove the decision to implement an internal crowdsourcing application. Some 529 employees from various levels, tenures, and locations across marketing, sales, finance, supply chain, market research, human resources, and research and development were invited to a website hosting the crowdsourcing tournament. To support an effective experiment test and control, current forecasting team members were not invited as participants. In order to maximize participation, we follow the well-known three wave data collection approach outlined by Dillman, Smyth, and Christian (2009). Of the sampling frame, 207 elected to register in the system with 154 making at least one prediction. We assess the self-selection bias present here to be no greater than any group of individuals that agree to complete a survey within an organization. There were no meaningful differences in respondents across the three waves.

The technology platform provides participants with a way to interface with the crowdsourcing application and other participants (Hopman, 2007). With a crowdsourcing prediction tournament (simply tournament hereafter), multiple participants indicate their predictions about future outcomes (e.g., projected sales increase, new product demand) on interactive continuous scales within the technology platform. The aggregated competitive predictions are presented as a distribution of predictions with a central value. Participants also use unique usernames in the technology interface to provide anonymity. This anonymity makes individuals more willing to join the tournament and share their own information and opinions as they are more likely to feel free from organizational influence and consequences (Hopman, 2007).

Common background information was made available when participants logged into the technology platform: historic sales trends, planned distribution, and promotional activities. Participants were encouraged to seek out and incorporate internal and external information beyond what was available to everyone. The aggregated collective prediction of all participants was posted publicly, in real-time, for each prediction question (see notes provided at the bottom of Fig. 1). After making their predictions, participants were able to view the aggregated predictions of all participants and update their own prediction.

The crowdsourcing tournament ran for six weeks in order to maximize opportunity for participation and to give the participants sufficient time to perform their own research and analysis.

The incentives built into the tournament provide motivation for participating, revealing good information, and performing well (Ho & Chen, 2007; Hopman, 2007). In each of the three divisions, the top two performers earned \$250 and \$150, respectively. Five additional \$50 prizes were awarded in each division through a lottery drawing based on total points earned. Points could be earned by predicting correct ranges, by indicating narrow confidence ranges, and by predicting early. Points were determined and prizes were awarded after actual outcome results became available (and could be compared to individual predictions).

For crowdsourcing applications to be successful, the outcomes to be predicted must be clearly defined and easily understood by participants (Bonabeau, 2009). For realism, predicted outcomes for this study were selected to represent a wide range of typical marketing decisions and forecasting methodologies used in practice: new product sales (Hardie, Fader, & Wisniewski, 1998), supply chain (Srivastava et al., 1999), and expected dollar sales (Rao & Bharadwaj, 2008). The questions in the field study consisted of questions the subject company currently forecasts and uses in their annual plan. Based on this design, a total of 11 questions are tested. For example, *What will be the U.S. unit sales for new product B for Q2 this year?* Each question has a carefully specified scale based on what is to be predicted (e.g., dollars, units), when (e.g., quarter, annual), and where (e.g., geography, channel). Participants used this same scale to indicate the range they thought the future actual outcome value would fall within. In order to avoid any anchoring bias that may occur, the current corresponding internal forecasts were not indicated on the scale or shared with the participants.

Participants were able to make forecasts for as many of the eleven questions they desired. They were able to return any number of times to update their forecasts or add forecasts up to the close of the tournament.

3.2. Comparable incumbent forecasts methodologies

The 11 crowdsourcing tournament predictions are compared directly against a corresponding set of incumbent company forecasting methods that varies widely by team, data, and methodology. For example, the market research function using judgmental consumer intention to buy survey as input into quantitative econometric causal model. Methods range from judgmental to quantitative as well as from basic extrapolation to complex modeling. All of the methods produced real forecasts, intended for regular business decisions and planning, for outcomes whose results would be tracked and compared against forecast.

3.3. Participant survey

After the tournament closed, an email invitation was sent to the 154 employees who had made at least one prediction. The email contained an embedded link to an online (Qualtrics) questionnaire, and noted that employees could win either a grand prize (an iPad) or one of several \$10 gift certificates for completing the survey. We received a total of 103 completed questionnaires, an overall response rate of 67%. Because we are studying the processes followed by the participants when they were making their individual predictions, it is necessary to sample from the original group of crowdsourcing participants. We assess the self-selection bias present here to be no greater than any group of individuals that agree to complete a survey within an organization. The demographic characteristics of the 103 survey respondents reveal no significant differences from the larger group of tournament participants.

3.4. Survey measures

The hypotheses state that the accuracy of predictions is driven by differences in *acquisition, distribution, and interpretation* factors, respectively. Accordingly, we developed reflective, multi-item scales for these three independent variables in three stages. First, we generate the initial measures by consulting the salient literature (e.g., market information processing, organizational theory, crowdsourcing, and collective intelligence). Second, a check of face and content validity was conducted through a panel of six senior marketing academics and two expert practitioners. Third, the survey instrument and scale items were pretested with 30 undergraduate students. Based on these exercises, we added some new scale items, and either modified or removed certain questions that were not clear.

Data + Collective Forecast

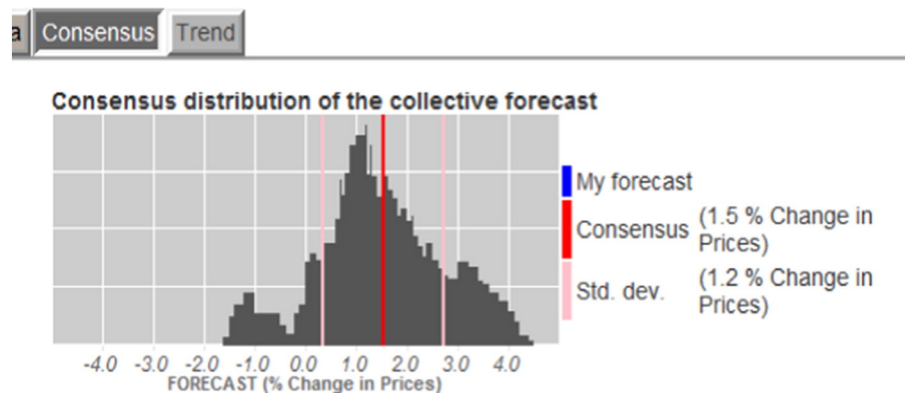


Fig. 1. Aggregated collective prediction display. Note: When entering into the prediction competition, participants go through the following steps: 1. Log into the competition technology platform. 2. Review system orientation and instructions. 3. Select a prediction question, and review background information provided for that question. (e.g., What will be the U.S. unit sales for new Product B for Q2 FY12?) This forecast is typically made by the brand marketing team using historical trend data on shipments and the marketing calendar. 4. Enter a prediction by setting the prediction range on the scale (similar to image) around the outcome which they expect. 5. Review aggregated collective prediction from the crowd (image above), and update their individual prediction as often as desired.

The scale development process yields 17 items (see Appendix A) deployed to measure the three crowdsourcing factors as latent constructs. Information acquisition is measured with a six-item scale that assesses the degree to which crowdsourcing participants seek additional information to use in their individual predictions; distribution is captured via a five-item scale assessing participant use of information shared in the tournament to make and update their own predictions; and interpretation is assessed via a six-item scale assessing differences in participant thinking and analysis while making their predictions.

In order to optimize the measurement model, individual scales were subjected to factor analysis and reliability analysis. Based on assessment guidance presented in Hair, Black, Babin, and Anderson (2010) and Kline (2011), seven poorly fitting scale items were removed. The following observations can be made regarding the remaining scale items: the cases-to-items ratio is 6:1, all scales have achieved unidimensionality, all KMO values are acceptable with the lowest value at .77, variance explained values range from 71.8% to 74.4%, all factor loadings exceed .80, all communality values exceed .60, all coefficient alphas exceed .90, and all item-total correlations exceed .70 (lowest item-total correlation exceeds .50).

To model the relationship between crowdsourcing factors and prediction accuracy, each participant was asked to select one prediction question they estimated in the tournament. They were then provided with a description of the team, data, and method for the corresponding internal forecast. For each crowdsourcing factor, respondents were asked to assess the difference between their prediction process (treatment condition) versus the internal forecast team and process (control) on seven point Likert-type scales (ranging from “1 = not at all” to “7 = a great deal”).

In the analysis the individual self-assessed, self-reported differences between prediction tournament participants and the internal forecast team (independent variables) are compared directly against individual prediction accuracy (APE) (dependent variable). Use of self-assessment scales, reference group comparisons, and difference measures are based on similar use in marketing (Li & Calantone, 1998; Maltz & Kohli, 1996) and organizational learning (March, 1991).

4. Analysis and results

4.1. Assessment of crowdsourcing prediction accuracy

The first research question asks whether crowdsourcing applications can improve market-oriented predictions within firms. This question needs to be answered before addressing the second research question regarding the mechanism underlying success. To answer the first research question, for each prediction task, we compare the prediction accuracy of (i) the aggregated collective predictions from the tournament, (ii) existing internal company forecasts, and (iii) actual post-tournament company results. The aggregated collective prediction is the simple average of the midpoints of all participants' final ranges submitted for each prediction question.

The assessment compares both tournament predictions and internal company forecasts against actual results for each of the 11 predictions. The tournament's predictions were available from the technology platform as soon as the tournament closed. The existing internal company forecasts were produced during regular company planning processes before the tournament was established. Actual results were reported through normal company reporting a few months after the tournament ended. Across 11 questions, 154 internal employees made 1460 predictions. Almost 70% of predictions were from participants from the R&D (31.2%), marketing (22.1%), and supply chain (16.9%) functions.

The most common measures of prediction or forecast accuracy are relative measures of accuracy provided by absolute percentage error (APE) and mean absolute percentage error (MAPE) (Armstrong, 2001; Armstrong & Collopy, 1992). APE in this analysis is calculated as the

difference between actual market results versus each of the tournament predictions and internal company forecasts:

$$APE_i = |\text{Actual}_i - \text{Prediction}_{ij}| / \text{Actual}_i \quad (1)$$

The subscript i represents prediction questions 1–11, and j represents the prediction methods: aggregated collective prediction (AC Pred) and internal forecast (Internal FCST). In this analysis, the MAPE is simply a mean of the set of APE values for particular predictions. The MAPE values for the aggregated tournament predictions and the internal company forecasts are 79.4% and 113.0%, respectively. A comparison of the MAPE scores indicates the mean values are significantly different from each other ($p = .0338$) with the prediction tournament displaying an average error 33.6% lower (95% CI [−.63, −.04]) than that of traditional forecasting methods employed by the subject firm.

As crowdsourcing applications often achieve performance statistically equal with existing internal forecast methods, the significant reduction in MAPE found here is one indication of improved results. Two additional comparison measures deemed to be reliable are relative reduction in error (RRE) and percent better (PB) (Armstrong, 2001; Armstrong & Collopy, 1992). A measure of relative reduction in error compares the reduction in error offered by an alternative prediction method compared to the traditional forecast method:

$$\text{AC Pred RRE}_i = \text{Internal FCST APE}_i - \text{AC Pred APE}_i \quad (2)$$

The subscript i represents prediction questions 1–11 and AC Pred RRE represents the relative reduction in error offered by the aggregated collective predictions compared to the internal company forecasts. The results are best interpreted through the percent better assessment which is the ratio of predictions that have a positive error reduction (RRE) to all predictions:

$$\text{AC Pred PB} = (\# \text{ of positive AC Pred RREs}) / (\text{Total} \# \text{ of Predictions}) \times 100 \quad (3)$$

Based on this calculation, the percent better ratio for the aggregated collective predictions is 72.7% (8 out of 11 questions). This result corroborates comparable analyses for Hewlett-Packard (Ho & Chen, 2007) and Intel (Hopman, 2007), which find a 75 percent better ratio. In addition to the prediction tournament reducing average error (MAPE) by almost 34 percentage points, it also reduces the error range by approximately 40% compared to the internal forecasts. It is interesting to note that areas where the aggregated collective predictions perform better are ones for which there was less company information available to support the prediction or internal forecast (e.g., new products, supply chain, and a new channel). These results answer research question one positively thus warranting the investigation into how crowdsourcing applications frequently produce more accurate predictions than traditional methods.

4.2. Examination of crowdsourcing factors

The 103 cases from the post survey and Smart-PLS 2.0 partial least squares modeling (PLS-SEM) are used to examine the second research question: how do crowdsourcing applications improve prediction accuracy? (Ringle, Wende, & Will, 2005). The sample size meets the minimum PLS-SEM threshold of 10 times the largest number of structural paths directed at a particular latent construct in the structural model (Hair, Ringle, & Sarstedt, 2011). The measurement model has also undergone rigorous scale analysis and purification to meet the requirement of a quality measurement model in order for PLS-SEM to yield acceptable parameter estimates with a smaller sample size (Hair, Sarstedt, Ringle, & Mena, 2012).

4.2.1. Measurement model assessment

From a confirmatory factor analysis with adjustments, the following observations can be made: composite reliability statistics range from .86 to .95 and AVE values range from .57 to .74. In addition, item loadings range from .51 to .97, no items cross-loaded onto other scales, and all items are significant through a bootstrapping procedure specified with no missing values, 5000 samples, and 103 cases (Hair et al., 2011; Hair et al., 2012).

Based on these results, it can be concluded that the measurement model meets or exceeds the requirements for reliability and validity in the PLS-SEM methodology.

4.2.2. Structural model assessment

Smart-PLS 2.0 was also run to analyze the effects that the acquisition, distribution, and interpretation factors (as latent exogenous constructs) have on prediction accuracy (as the measured dependent variable). All values for Stone–Geisser's Q2 are greater than zero indicating that explanatory constructs display adequate predictive relevance and validity to predict their indicators (Hair et al., 2011; Hair et al., 2012). The coefficient of determination (R-square), the primary criterion for structural model assessment in PLS-SEM, is .103 for this model (Hair et al., 2011; Hair et al., 2012).

The standardized path coefficients and t-statistics are produced through a standard bootstrapping procedure specified with no missing values, 5000 samples, and 103 cases (Hair et al., 2011; Hair et al., 2012). Based on these scales and tests, the null hypothesis should be rejected for H1 Acquisition (path coefficient $-.238$, single tailed $t = 2.149$) and H3 Interpretation (path coefficient $-.164$, $t = 1.514$) indicating that higher levels of information acquisition and differences in interpretation both lead to improved prediction accuracy (as decreased APE) in a crowdsourcing application. Conversely, the null hypothesis should fail to be rejected for H2 Distribution ($t = .456$) based on a confidence level of 90%, which is appropriate for the purposes of an exploratory examination of drivers in forecasting, especially with a smaller sample (Gartner & Thomas, 1993).

The lower R-square value (.103) can be partially explained by strains on internal validity due to heavy influence of extraneous variables related to a field study (Churchill & Iacobucci, 2005) and to a significant number of factors that influence the accuracy of market demand forecasts (Gartner & Thomas, 1993; Kahn, 2002). For example, Gartner and Thomas (1993) identify 30 factors that influence forecasting accuracy. Considering the massive amounts of wasted human effort and money that can result from errors, being able to explain even 10% of the variation in forecasting outcomes could have a significant impact on decisions and financial outcomes (Kahn, 2002).

This research establishes that accuracy of a crowdsourcing application is driven most by (i) information acquisition (i.e., participants obtaining additional information to use in their individual predictions) and (ii) differences in interpretation (i.e., the extent to which participants bring different perspectives and mental models to their prediction). The findings for these two hypotheses (H1 and H3) correspond with theory presented in the supporting literatures. For example, March (1991) finds that new information brought into the organization and learning heterogeneity are two of the most important determinants of organizational learning.

We believe that the lack of support for distribution (H2) may be due to most participants in the field study accessing and using the available information, and thereby limiting variance in this scale. The design of the technology interface and initial instructions likely increased awareness of the shared information available in the tournament, and may have led participants to believe its use was a necessary part of the process. Future studies might create designs with varied levels of orientation of shared information beyond basic awareness (e.g., not instructing and encouraging use in one condition).

5. Discussion

There has been limited empirical testing and, more importantly, virtually no explanation as to how crowdsourcing platforms produce superior results. This article thereby seeks to evaluate: can crowdsourcing applications improve market information processing and resulting predictions within firms and, if so, how?

5.1. Theoretical contributions

Regarding the first research questions, this article advances theory in three ways. First, it responds directly to a specific knowledge gap identified in the forecasting literature to compare the efficacy of crowdsourcing applications to traditional forecasting approaches in realistic marketing settings (Armstrong, 2006; Spann & Skiera, 2003).

Second, this study addresses several questions regarding how firms can improve their capabilities through improving their development and use of market knowledge. Specifically, this research demonstrates a method to better integrate new and existing market knowledge (Vorhies et al., 2011), satisfies the need to develop new tools for internal analysis (Morgan, 2012), and introduces a new market-based learning tool that can enhance market information management (Vorhies & Morgan, 2005).

Third, to the best of our knowledge, this study is the first empirical test of crowdsourcing's efficacy to support real decisions within a marketing context. It is also the first application of non-trading-based crowdsourcing to business forecasting. From our test and analysis, we find crowdsourcing provides more accurate predictions than traditional forecasting methods in almost three-fourths of the cases. In addition, the application reduces average error rates by over 33% and error ranges by 40%. This result answers the first research question and corroborates comparable analyses for Hewlett-Packard (Ho & Chen, 2007) and Intel (Hopman, 2007) which find the same ratio.

To evaluate the second research question (i.e., how does crowdsourcing actually produce better results?), we go beyond the mere replication crowdsourcing performance from other disciplines and make three additional contributions. First, our most important finding is that differences in how a crowdsourcing application performs acquisition and interpretation of market information leads to superior forecasts of market-oriented outcomes. The acquisition of more and different market information (H1) and differences in interpretation of the participants (H3) are the mechanisms that drive greater accuracy of crowdsourcing applications. These results corroborate the earlier suggestions that information heterogeneity and new information brought into the firm are two of the most important contributors to organizational learning (March, 1991; Sinkula, 1994; Slater & Narver, 1995). Similarly, they provide support for the conjecture advanced in the writings on crowdsourcing and collective intelligence that diversity of participants is an important driver of prediction accuracy (Surowiecki, 2005). These findings respond directly to criticism in the literature that while published studies have shown that collective intelligence tools tend to win the horse race against competing methodologies, the underlying mechanisms that produce superior results have not been evaluated (Green et al., 2007).

Second, this study makes a contribution to understanding the ideal use and boundaries of crowdsourcing (Green et al., 2007). It finds that crowdsourcing performs better for outcomes where there is limited information or historic data to support predictions, and where there may be greater levels of uncertainty. In the current study, these are innovations in products, supply chains, and entering new channels. Conversely, in areas where there is more experience and abundant historical data (overall sales predictions), crowdsourcing may offer less or no improvement in accuracy compared to the incumbent methodologies. A plausible reason for this finding is that typical forecasting statistical models such as ARIMA (autoregressive integrated moving average)

or exponential smoothing are able to use large series of historical data to calibrate tools and interpret results effectively.

Third, these findings contribute to how market information processing can be improved by adopting new and innovative decision support tools from other disciplines. The introduction of crowdsourcing into marketing, which requires drawing upon several literatures outside of marketing, responds directly to the MSI's research priorities calling to expand marketing's academic boundaries to leverage alternative forms of information and unstructured data to develop better, real-time intelligent systems. This study also responds to the call to better understand the benefits and use of decision support tools that (in spite of their benefits) are under-utilized in marketing (Day, 2011; Krishnan & Ulrich, 2001). The uniquely adapted way that crowdsourcing applications perform market information processing applied to specific marketing questions responds to calls in the literature to find new and innovative tools and techniques to support the development and use of market learning (Day, 2011).

There are two additional areas where crowdsourcing may further the understanding of market information processing and its outcomes. Huber (1991); Moorman (1995), and Sinkula (1994) all discuss how market learning occurs at both the individual and organizational levels, but raise questions regarding how individual level processing becomes organizational, and whether there are other intermediate levels that may be useful. For example, an issue that mitigates knowledge integration in organizations is that the majority of the collective intelligence is tacit as it resides in employees' minds; therefore, it cannot be readily codified and passed along to others (Moorman, 1995). As a result, scholars have been guided to investigate new ways to tap into dispersed tacit knowledge so that valuable experiences, insights, and expertise can be widely shared and utilized (Day, 2011; Day & Nedungadi, 1994). Our results demonstrate that crowdsourcing platforms are powerful tools to solicit input from employees regarding a certain task, aggregate the diverse pieces of information scattered throughout the firm, and harness it to yield superior predictions. These results respond to the problems raised by Huber (1991) and Sinkula (1994) about where information resides or is stored, how it is retrieved for use, and the role of technology in facilitating this connection and application.

This research also contributes to the noted need in the marketing–finance interface literature to move beyond point estimates in forecasting sales (Rao & Bharadwaj, 2008). Crowdsourcing can combine a diverse set of individual revenue predictions into one estimate with a probability distribution, thereby presenting a useful mechanism for estimating both the revenues and uncertainty needed to link marketing recommendations and actions to expected cash flows.

5.2. Managerial contributions

The current study makes several valuable managerial contributions. First, the improvement in accuracy demonstrated here introduces crowdsourcing as a valuable new decision support tool for consideration by business managers. The high external validity of the current study's findings (actual employees using real data to solve live problems) provides strong evidence that crowdsourcing can yield market-oriented predictions equal to or better than those produced by the incumbent methods used by the firm. Crowdsourcing applications are also less expensive, and faster and easier to implement than many existing methods (e.g., consumer surveys) (Ho & Chen, 2007; Spann & Skiera, 2003).

Second, from the field implementation, we learned that executives at our host company were not convinced by studies that crowdsourcing can generate improved predictions. Rather, they demand insight into why they work. This fact is evidenced in a comment made by an executive during our initial meetings: "Unless I can explain how these things work, I cannot introduce them to my management. I cannot ask them to accept numbers from a magic black box." In direct response to this challenge, the current study demonstrates and explains how

crowdsourcing actually produces superior results. This research shows managers that they can achieve improved predictions by encouraging participants to collect additional information for their predictions (H1). Executive teams can encourage participants to "do their homework," which this study shows can improve the accuracy of the aggregated predictions. The results also show that designing a crowdsourcing application that fosters heterogeneity improves individual predictions. Thus, executives should seek to gain as many varied perspectives as possible when implementing crowdsourcing so that unique pieces of information possessed by disparate people can make their way into the collective organizational consciousness (H3).

Third, managers can use crowdsourcing to support multiple decisions such as evaluating new markets and channels to enter, selecting the best new concepts or ventures to pursue, developing new merchandising and promotional programs, and investigating the impact of changes to pricing structures (Gartner & Thomas, 1993; Hardie et al., 1998; Kindström, Kowalkowski, & Sandberg, 2013).

Fourth, it is important to note that crowdsourcing can be used to support many business decisions outside of the marketing discipline. For example, as examined in this study, analysis and planning for supply chain, finance, and annual operating planning (Rao & Bharadwaj, 2008; Srivastava et al., 1999). Crowdsourcing can also be used to assess trends and changes in marketplace and environmental factors that have a significant impact on their operations and profitability: predictions on future crop and fuel prices, general category spending behaviors, and other macro industry and economic variables.

5.3. Limitations and future research directions

The participants for our field test were employees internal to the subject firm. Thus, parties outside of the organization (channel partners, suppliers, advertising agencies, etc.) were not represented. The importance of information acquisition and differences in interpretation in improving prediction accuracy demonstrated in this study suggests that it would be impactful to include individuals external to the firm in crowdsourcing.

The prediction competition that consisted of 11 questions did outperform the traditional forecasting methodologies the majority of the time; however, there were instances in which the incumbent approaches were better at estimating future outcomes. It would be revealing in future research to delve further into the conditions under which crowdsourcing yields superior versus sub-par estimates.

The small sample size of the post survey ($n = 103$) likely affected the statistical tests. Although sufficient for PLS-SEM analysis, it is possible that the third factor (shared information) might have been found significant with a slightly larger sample. Future researchers can benefit from our prediction test for guidance. We found that 29% of the 529 employees invited made at least one prediction in Study 1 (i.e., $n = 154$). If the goal of future research is to simply compare predictions generated by a crowdsourcing to those stemming from other methodologies, then recruiting approximately 100 employees should suffice. Earlier writings have noted that securing as few as 30 individuals (Ho & Chen, 2007; Servan-Schreiber, 2012, 2013) from different functional areas and hierarchical levels (Surowiecki, 2005) will permit comparisons. If, however, a researcher wishes to further explore the factors that can lead to improved accuracy, our conversion ratio would suggest that recruiting 1000 employees can yield a sample size closer to 200. We did not analyze variances in forecast accuracy across organizational demographic information (marketing, R&D, etc.). Future studies with larger sample sizes could investigate how individual differences could influence forecast accuracy.

We made a methodological decision to use a non-trading crowdsourcing method over a trading-based method in order to simplify the experience for participants and improve participation and engagement. We based this decision on literature finding that trading-based platforms encumbered with trading complexity that can make

them less attractive for large lay audiences. Non-trading-based platforms have been found easier for people with no stock trading experience to use and are to be more suitable for use in corporate settings (Servan-Schreiber, 2013). Testing the actual differences between trading versus non-trading platforms was not one of our objectives and would be an interesting area of investigation.

In closing, we reiterate that there has been limited empirical testing and, more importantly, virtually no explanation as to how crowdsourcing platforms produce superior results. This article takes a forward step towards evaluating crowdsourcing, with the hope that future research can attend to the open questions that remain.

Appendix A. Constructs and scale items

Acquisition In making your prediction, to what degree did you...

(1 = Not at all...7 = A great extent)

1. Try to search for *more* information similar to the description (see note below)?
2. Try to search for *different* information than in the description?
3. Eventually obtain more information similar to the description?
4. Eventually obtain different information than in the description?
5. Collect information from sources outside your business unit?
6. Look for information about customers or competitors?

Distribution

In making your prediction, to what degree did you...

(1 = Not at all...7 = A great extent)

Look at the consensus prediction available in the system? (the displayed...)

1. Look at the combined prediction available in the system?
2. Find the shared supporting information adjusted your thinking?
3. Find the consensus prediction in the system adjusted your thinking?
4. Find the shared supporting information useful?
5. Find the consensus prediction in the system useful?

Interpretation

Compared to the above description, to what degree did you...

(1 = Very similar...7 = Very different)

1. Have a different perspective than the people in the description? (e.g., viewpoint)
2. Have different knowledge than the people in the description? (e.g., local, specialized)
3. Have different experience than the people in the description?
4. Have different skills and abilities than the people in the description?
5. Have different tools and techniques than the people in the description?
6. Think your work area has different resources than the people in the description?

Note: The references above to “the description” pertain to internal forecast method descriptions participants are directed to within the survey. Corresponding references for scale items are available upon request.

References

- Afuah, A., & Tucci, C. L. (2012). Crowdsourcing as a solution to distant search. *Academy of Management Review*, 37(3), 355–375.
- Armstrong, S. J. (2001). *Principles of forecasting: A handbook for researchers and practitioners*. Norwell, MA: Kluwer Academic Publishers.
- Armstrong, S. J. (2006). Findings from evidence-based forecasting: Methods for reducing forecast error. *International Journal of Forecasting*, 22(3), 583–598.
- Armstrong, S. J., & Collopy, F. (1992). Error measures for generalizing about forecasting methods: Empirical comparisons. *International Journal of Forecasting*, 8, 69–80.
- Arrow, K. J., Forsythe, R., Gorham, M., Hahn, R., Hanson, R., Ledyard, J. O., ... Zitzewitz, E. (2008). The promise of prediction markets. *Science*, 320(5878), 877–878.
- Bharadwaj, N., & Noble, C. H. (2015). Innovation in data-rich environments. *Journal of Product Innovation Management*, 32(3), 476–478.
- Bonabeau, E. (2009). Decisions 2.0: The power of collective intelligence. *Sloan Management Review*, 50(2), 45–52.
- Cepeda, G., & Vera, D. (2007). Dynamic capabilities and operational capabilities: A knowledge management perspective. *Journal of Business Research*, 60(5), 426–437.
- Churchill, G. A., & Iacobucci, D. (2005). *Marketing research: Methodological foundations* (9th ed.). Thompson South-Western.
- Cowgill, B., Wolfers, J., & Zitzewitz, E. (2009). Using prediction markets to track information flows: Evidence from Google. *Working paper, Dartmouth University*.
- D'Aveni, Richard A., Dagnino, Giovanni B., & Smith, Ken G. (2010). The age of temporary advantage. *Strategic Management Journal*, 31(3), 1371–1385.
- Day, G. S. (1994). The capabilities of market-driven organizations. *Journal of Marketing*, 58(4), 37–52.
- Day, G. S. (2011). Closing the marketing capabilities gap. *Journal of Marketing*, 75(4), 183–195.
- Day, G. S., & Nedungadi, P. (1994). Managerial representations of competitive advantage. *Journal of Marketing*, 58(2), 31–44.
- Diamantopoulos, A., & Winklhofer, H. (2003). Export sales forecasting by UK firms: Technique utilization and impact on forecast accuracy. *Journal of Business Research*, 56(1), 45.
- Dillman, D. A., Smyth, J. D., & Christian, L. M. (2009). *Internet, mail, and mixed-mode surveys: The tailored design method* (3rd ed.). New Jersey: John Wiley & Sons.
- Eisenhardt, K. M., & Martin, J. A. (2000). Dynamic capabilities: What are they? *Strategic Management Journal*, 21(10/11), 1105.
- Gartner, W. B., & Thomas, R. J. (1993). Factors affecting new product forecasting accuracy in new firms. *Journal of Product Innovation Management*, 10(1), 35–52.
- Good Judgment Project (2014). Harnessing the wisdom of the crowd to forecast world events. (accessed January 15, 2014, available at) <https://www.goodjudgmentproject.com/>
- Green, K. C., Armstrong, J. S., & Graefe, A. (2007). Methods to elicit forecasts from groups: Delphi and prediction markets compared. *Foresight: The International Journal of Applied Forecasting*, 8(8), 17–21.
- Hair, J. F., Black, W. C., Babin, B. J., & Anderson, R. E. (2010). *Multivariate data analysis* (7th ed.). New Jersey: Prentice Hall.
- Hair, J. F., Ringle, C. M., & Sarstedt, M. (2011). PLS-SEM: Indeed a silver bullet. *Journal of Marketing Theory and Practice*, 19(2), 139–151.
- Hair, J. F., Sarstedt, M., Ringle, C. M., & Mena, J. A. (2012). An assessment of the use of partial least squares structural equation modeling in marketing research. *Journal of the Academy of Marketing Science*, 40(3), 414–433.
- Hardie, B., Fader, P. S., & Wisniewski, M. (1998). An empirical comparison of new product trial forecasting models. *Journal of Forecasting*, 17(34), 209–229.
- Ho, T. H., & Chen, K. Y. (2007). New product blockbusters: The magic and science of prediction markets. *California Management Review*, 50(1), 144–158.
- Hopman, J. W. (2007). Using forecasting markets to manage demand risk. *Intel Technology Journal*, 11(2), 127–135.
- Horn, C. F., Ohneberg, M., Ivens, B. S., & Brem, A. (2014). Prediction markets—A literature review 2014 following Tziralis and Tsiopoulou. *Journal of Prediction Markets*, 8(2), 89–126.
- Huber, G. P. (1991). Organizational learning: The contributing processes and the literature. *Organization Science*, 2(1), 88–115.
- Kahn, K. B. (2002). An exploratory investigation of new product forecasting practices. *Journal of Product Innovation Management*, 19(2), 133–143.
- Kindström, D., Kowalkowski, C., & Sandberg, E. (2013). Enabling service innovation: A dynamic capabilities approach. *Journal of Business Research*, 66(8), 1063–1073.
- Kline, R. B. (2011). *Principles and practice of structural equation modeling*. New York: The Guilford Press.
- Krishnan, V., & Ulrich, K. T. (2001). Product development decisions: A review of the literature. *Management Science*, 47(1), 1–21.
- Li, T., & Calantone, R. J. (1998). The impact of market knowledge competence on new product advantage: Conceptualization and empirical examination. *Journal of Marketing*, 62(4), 13–29.
- Malone, T. W., Laubacher, R., & Dellarocas, C. (2010). Harnessing crowds: Mapping the genome of collective intelligence. *Sloan Management Review*, 51(3), 21–31.
- Maltz, E., & Kohli, A. K. (1996). Market intelligence dissemination across functional boundaries. *Journal of Marketing Research*, 33(1), 47–61.
- March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, 2(1), 71–87.
- Marketing Science Institute (2014). *2014–2016 Research priorities*. Cambridge, MA: MSI.
- Matzler, K., Grabher, C., Huber, J., & Füller, J. (2013). Predicting new product success with prediction markets in online communities. *R&D Management*, 43(5), 420–432.
- Milliken, F. J., & Martins, L. L. (1996). Searching for common threads: Understanding the multiple effects of diversity in organizational groups. *Academy of Management Review*, 21(2), 402–433.
- Moorman, C. (1995). Organizational market information processes: Cultural antecedents and new product outcomes. *Journal of Marketing Research*, 32(3), 318–335.
- Morgan, N. A. (2012). Marketing and business performance. *Journal of the Academy of Marketing Science*, 40(1), 102–119.
- Narver, J. C., & Slater, S. F. (1990). The effect of a market orientation on business profitability. *Journal of Marketing*, 54(October), 20–35.
- Neff, J. (2012, March 26). Is Procter & Gamble losing its edge? *Advertising Age*.
- Rao, R. K. S., & Bharadwaj, N. (2008). Marketing initiatives, expected cash flows, and shareholders' wealth. *Journal of Marketing*, 72(1), 16–26.
- Ringle, C. M., Wende, S., & Will, A. (2005). *SmartPLS 2.0 (beta)*. taken from <http://www.smartpls.de> Hamburg, Germany: SmartPLS.
- Rodan, S., & Galunic, C. (2004). More than network structure: How knowledge heterogeneity influences managerial performance and innovativeness. *Strategic Management Journal*, 25(6), 541–562.
- Servan-Schreiber, E. (2012). Prediction markets: Trading uncertainty for collective wisdom. In H. Landemore, & J. Elster (Eds.), *Collective wisdom: Principles and mechanisms*. New York: Cambridge University Press.
- Servan-Schreiber, E. (2013). Trading uncertainty for collective marketing wisdom. *Presentation at 2013 American Marketing Association Winter Educators' Conference*.
- Sinkula, J. M. (1994). Market information processing and organizational learning. *Journal of Marketing*, 58(1), 35–45.
- Sinkula, J. M., Baker, W. E., & Noordewier, T. (1997). A framework for market-based organizational learning: Linking values, knowledge, and behavior. *Journal of the Academy of Marketing Science*, 25(4), 305–318.
- Slater, S. F., & Narver, J. C. (1995). Market orientation and the learning organization. *Journal of Marketing*, 59(3), 63–74.
- Slater, S. F., & Narver, J. C. (2000). Intelligence generation and superior customer value. *Journal of the Academy of Marketing Science*, 28(1), 120–127.

- Soukhoroukova, A., Spann, M., & Skiera, B. (2012). Sourcing, filtering, and evaluating new product ideas: An empirical exploration of the performance of idea markets. *Journal of Product Innovation Management*, 29(1), 100–112.
- Spann, M., & Skiera, B. (2003). Internet-based virtual stock markets for business forecasting. *Management Science*, 49(10), 1310–1326.
- Srivastava, R. K., Shervani, T. A., & Fahey, L. (1999). Marketing, business processes, and shareholder value: An organizationally embedded view of marketing activities and the discipline of marketing. *Journal of Marketing*, 63(4), 168–179.
- Surowiecki, J. (2005). *The wisdom of crowds*. New York: Anchor Books.
- Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7), 509–533.
- Vorhies, D. W., & Morgan, N. A. (2005). Benchmarking marketing capabilities for sustained competitive advantage. *Journal of Marketing*, 69(1), 80–94.
- Vorhies, D. W., Orr, L. M., & Bush, V. D. (2011). Improving customer-focused marketing capabilities and firm financial performance via exploration and exploitation. *Journal of the Academy of Marketing Science*, 389(1), 736–756.
- Wierenga, B., Van Bruggen, G. H., & Staelin, R. (1999). The success of marketing management support systems. *Marketing Science*, 18(3), 196–207.
- Yan, R., & Ghose, S. (2010). Forecast information and traditional retailer performance in a dual-channel competitive market. *Journal of Business Research*, 63(1), 77–83.