From Online Content to Offline Results: A Field Experiment Examining the Effects of an Enterprise Social Network "Best Practices" Initiative

Shelley Xin Li
Leventhal School of Accounting
University of Southern California, Los Angeles, CA 90089

Tatiana Sandino
Harvard Business School
Harvard University, Boston, MA 02163

April 16, 2021

Abstract

Enterprise social networks (ESNs) supplement traditional top-down knowledge-sharing channels within organizations with knowledge from a broad set of employees. This enables participants facing different circumstances to benefit from a wide range of ideas (Nonaka and von Krogh 2009). However, the large volume and varied quality of user (employee)-generated contents can make it difficult to find relevant and trustworthy information to achieve better performance. We conducted a natural field experiment in collaboration with a large grocery store chain in Europe that introduced a "best practices initiative" on its ESN encouraging high-performing units to share business practices. This intervention shifted more of the employees' activities (i.e., posts and comments) to the ESN groups where the best practices were introduced. We do not find a significant, immediate improvement in sales. However, our results show a significant change on the sales trend for the treatment stores implementing the initiative relative to the control stores (resulting in 3.67-percent higher sales over the course of the four-month initiative). These results suggest that best practices initiatives on ESNs lead to performance improvement over time consistent with gradual learning from the best practices posts. This sales trend effect was driven by the regions showing greater activity related to their "best practices" posts (i.e., more employees seeing and reacting to the posts on the ESN). Furthermore, the effects of the initiative were more positive in stores with lower ex-ante exposure to information on business practices from peer stores, in stores serving markets that were more similar to those of the best-practices ("sharing") units, in stores that had lower performance prior to the intervention, and in regions where employees had lower (rather than higher) trust in their regional managers' competence before the intervention. Overall, these results suggest that systematically sharing contents from high-performing units on an ESN could focus employees' attention to productive knowledge sharing on an ESN and improve sales, especially under the aforementioned conditions.

The researchers obtained human subjects research approval from their universities to conduct this study. Furthermore, the participating universities signed a data usage agreement with the participating company, compliant with the General Data Protection Regulation of the European Union. The agreement specifies the research team members who can access the data. We thank Evelina Forker and seminar participants at Arizona State University, Carnegie Mellon, and the University of Southern California, as well as participants at the 2020 Consortium for Operational Excellence in Retailing at Harvard Business School and the 2021 Management Accounting Section of the American Accounting Association's Midyear Meeting for helpful comments. We thank Lene Furuseth, Victoria Liublinska Prince, Trang Nguyen, and Christine Rivera for excellent research assistance throughout the project. We also thank the grocery store chain company for agreeing to collaborate with us on this field experiment, supporting us with the development of this project, and providing us with access to data. All errors are our own.

I. INTRODUCTION

Companies have increasingly adopted enterprise social networks (hereafter, ESNs)—digital information systems enabling employees to share information with, socially interact with, and observe information exchanges among other employees. Employees post content (pictures, videos, texts), share posts, "like" posts, and/or engage in dialogues about posted ideas through ESNs. A 2019 survey by Gartner reported that 45 percent of employees relying on digital technologies across a variety of companies in the US, Europe, and the Asia-Pacific countries used ESNs daily.[1] More recently, publicly traded companies such as Microsoft (owner of two ESNs, Teams and Yammer), Slack Technologies, and Facebook (owner of the ESN Workplace) have reported that ESN adoption has skyrocketed throughout the world, following the quarantine measures adopted in 2020 in response to the COVID-19 pandemic.[2]

A chief reason for implementing ESNs is to promote knowledge sharing—that is, to encourage employees from different parts of the organization to communicate with each other and share ideas on how to do their work—as a mechanism to drive performance improvements. Several studies have shown the potential of ESNs to promote knowledge sharing by enabling employees to learn "who knows what" and "who knows whom" (Leonardi [2018], Neeley and Leonardi [2018]). This knowledge enables participants to search for information relevant to their particular circumstances. However, some studies claim that ESNs generally fail to promote effective knowledge sharing or

[1] Between March and April of 2019, the Gartner Digital Worker Survey surveyed 7,261 full-time employees who used digital technologies for work purposes in companies with 100 workers or more (Costello 2019).

[2] For example, in its April 29, 2020, conference call presentation, Microsoft executives indicated that the number of organizations adopting its Teams app had tripled in the last two months and that Teams had grown to 44 million users conducting over 2.7 billion minutes of online meetings daily. Slack Technologies, in its first-quarter earnings call on June 4, 2020, indicated that revenue was up 50% year-on-year and that it had added a record 90,000 new organizations that quarter, bringing its total to over 750,000. They added: "For our paid users, average time spent actively using Slack each day increased from just under 90 minutes at the end of Q4 to over 120 minutes per day at the end of Q1."

improve performance due to (a) opportunity costs and risks incurred by employees posting information (Levine and Prietula [2012]) (b) the search costs for consuming information (Levine and Prietula [2012]), and (c) lack of leadership and guidance on how to navigate the ESN (Charki, Boukef, and Harrison [2018]). In order to limit these downside costs and to enhance productive knowledge sharing in ESNs, scholars and practitioners have recommended best practices initiatives that share insights on how high-performing units conduct business with other units via ESNs (Oleson [2013], Charki, Boukef, and Harrison [2018]).

In this paper, we seek to understand whether and when an initiative promoting best practices from high-performance units (hereafter, best practices or BP units) on an ESN could effectively guide employees' online attention, leading to better offline performance. Employees in best practices units are likely to have an idea of which actions drive results. Directing attention to these units' best ideas could enhance learning by reducing the cost to other units (hereafter, learning units) of searching for good ideas and processing information. But whether best practices initiatives enhance knowledge sharing and performance is unclear. The effectiveness of such initiatives may depend on the relevance and novelty of the ideas shared by the best-practices units and the extent to which employees are willing and able to consume those ideas. Such initiatives could also backfire if they discourage employees (other than those in best practices units) from sharing content and/or if they limit the breadth of ideas employees pay attention to on an ESN (that is, one of the key advantages of sharing knowledge horizontally, via a social network, rather than vertically, from the top down).

We examine the effects of an initiative promoting best practices through an ESN at a large grocery store chain in Europe whose stores were operated by franchisees. Prior to the intervention, the company shared best practices relying on top-down methods (e.g., operating guidelines and

training workshops). It also used an ESN similar to Facebook that enabled employees to develop groups and engage in discussions as they saw fit. Among its many groups, the ESN included one "regional group" for each sales region, accessible by all employees of stores in that region. The best practices initiative was conducted as a natural field experiment that involved 587 stores across 11 regions. The treatment group included 274 stores in five randomly selected regions. The control group included 313 stores in six control regions. The treatment consisted of promoting over the ESN—every month, in each region—best practices from two high-performing best practices stores in that region. To makes sure that the initiative was implemented in a consistent manner across the treatment regions, twice a month, the HR team asked a regional sales manager (an internal "reporter") to interview a manager from a high-performing store and post an album on the ESN presenting best practices from that store, accessible to all store employees in that region (and in that region only). The albums were posted following the same standardized format across all regions. The initiative was neither implemented nor discussed in the control regions.

We conduct a difference-in-differences analysis to test the main effects of the initiative. We find that the initiative immediately led employees to increase their activities (posts and comments) in the regional groups where the best practices were introduced. However, we do not observe an immediate improvement in sales. Instead, we find that the initiative was associated with an improvement in sales trends in the treatment group relative to the control group. This change in sales trend resulted in 3.67 percent higher sales in the treatment group than in the control group by the end of the 18-week initiative. We find that the effects of the initiative were stronger in regions where the posts gathered greater attention, consistent with the expectation that the posts drove the increase in sales. Furthermore, we find stronger effects in stores with *less* preexisting knowledge of best practices (measured based on the store's lack of geographic exposure to other same-

company stores), in stores that served markets geographically closer to the best-practices units (that is, in stores where the knowledge shared was most likely to be relevant), in previously lower-performing stores (that is, those with below median sales-per-employee before the intervention), and in regions where employees perceived their regional managers to be *less* competent (suggesting the initiative substituted for lack of advice from management). Overall, these results suggest that a best practices initiative on an ESN could improve sales, especially under the aforementioned circumstances.

We contribute to the literature in accounting that studies the effect of management control systems on learning and knowledge sharing. With a few exceptions, this research has focused on the effects of incentives and monitoring. Regarding incentives, prior studies show that employees are more willing to share knowledge with their peers when their compensation incorporates subjectivity (Cheng and Coyte [2014]) and that companies that value knowledge sharing are more likely to tie incentives to output (rather than input) measures and to group (rather than individual) measures (Hwang, Erkens, and Evans [2009]). Regarding monitoring, Campbell, Epstein, and Martinez-Jerez [2011] find that loosely monitored business units show higher learning rates than tightly monitored ones. Our study extends this literature by examining the performance and learning effects of sharing best practices through the kind of digital information system that have been increasingly adopted by companies: an ESN that aims to facilitate learning and performance. Our findings suggest that such sharing of best practices increases the rate of performance improvement. Our study also contributes to an emerging literature in accounting, information systems, and management that examines the effects of information-sharing systems on learning and performance. This research has shown mixed results. While some cross-sectional analyses show positive results of information-sharing systems on performance and innovation (e.g., Kuegler,

Smolnik, and Kane [2015], Leonardi [2018]), other studies examining causal relations more closely have found such positive effects only under certain circumstances and negative effects under other circumstances (Levine and Prietula [2012], Li and Sandino [2018]). We examine the impact of a best practices initiative as a mechanism used on ESN to simultaneously reduce employees' search costs and accelerate knowledge acquisition. We evaluate whether such an initiative can improve knowledge sharing and, if so, the circumstances under which it is most likely to do so.

Finally, our study sheds light on ways to leverage ESNs to increase productive employee interactions and exchanges of ideas, a topic which is of great interest to business. To the best of our knowledge, our study is the first to examine whether a best practices initiative on an ESN can enhance idea sharing across units and improve financial performance.

The rest of the paper is organized as follows: Section II presents our hypothesis development. Section III describes the research setting and method. Section IV explains our research design, presents our analyses, and discusses the findings. Section V concludes.

II. HYPOTHESIS DEVELOPMENT

2.1. What Are the Effects of a Best Practices Initiative on an ESN?

An emerging body of literature in accounting and management suggests that ESNs can benefit organizations. A study of a financial services organization suggests that an ESN system can lead to greater innovation and less duplication of work among participants who learn "vicariously rather than through experience" (Leonardi [2014]: 796). Using data from another financial institution, Campbell, Erkens, and Loumioti [2014] find that lending officers sharing more idiosyncratic

information through exception reports on an information-sharing system (documenting arguments used by loan officers making loan exceptions) granted more loans and reduced both interest rates and loan charge-offs. Based on a cross-sectional survey, Kuegler, Smolnik, and Kane [2015] find that ESNs are associated with improved task performance when employees share information within their teams and with greater innovation when employees share information across teams. Finally, in a mobile phone retail chain that introduced an information-sharing system on which employees shared marketing posters they had created, Li and Sandino [2018] find that the system was associated with an increase in creativity (but not in financial performance) in stores using the system more frequently.

Despite the evidence above suggesting that an ESN can lead to productive knowledge exchanges, those studies do not show an overall association between ESN and financial performance. Furthermore, various studies contend that a large percentage of companies attempting to leverage or sustain productive knowledge sharing on ESNs fail to do so (Charki, Boukef, and Harrison [2018], Neeley and Leonardi [2018]). Some studies highlight that the costs associated with interpersonal exchanges often exceed the benefits of sharing information (Levine and Prietula [2012]). There are multiple reasons for this. Some employees are reluctant to share their best ideas, especially when they perceive their opportunity costs to be high and/or fear giving away relative advantages in situations in which business units compete for scarce resources and/or opportunities (Butt, Antia, Murtha, and Kashyap [2018], Li and Sandino [2018]). Others hold back their questions or opinions for fear of being criticized or seen as spending too much time on the ESN (Neeley and Leonardi [2018]). Still others avoid sharing or seeking knowledge because they do not consider the other ESN participants trustworthy (Neeley and Leonardi [2018]). Even in organizations whose employees feel comfortable sharing knowledge, searching for relevant ideas

through unstructured ESN interactions—potentially overloaded with information irrelevant to a particular individual—can be time-consuming and/or unproductive (Levine and Prietula [2012]).

One of the initiatives that scholars and practitioners have recommended for enhancing productive knowledge sharing on ESNs is to regularly feature the practices of high-performing or exemplary units (or employees) as best practices (Oleson [2013], Charki, Boukef, and Harrison [2018]). We define best practices initiatives implemented on ESNs as those consisting of sharing actions considered effective at high-performing units with other units via ESNs. We expect such best practices initiatives to improve performance because they can (a) motivate productive knowledge sharing on the part of high-performing units (otherwise reluctant to share their knowledge on an ESN) by formally recognizing their distinctive practices and ideas and/or creating a mechanism through which they can benefit the organization;[3] (b) uncover valuable ideas that learning units could use to improve performance; (c) legitimize ideas and content shared by participants by presenting the featured best practices from widely trusted sources in a compact, organized format; (d) increase overall engagement through the discussion of more focused ideas relevant to the units' work; and (e) minimize the costs to ESN users of searching for relevant information.

Although a best practices initiative on an ESN could promote knowledge sharing, it is not obvious that it would improve financial performance. Employees in high-performing units may not know what actions led to their success and, consequently, may not be able to share them. Prior literature also warns that internal competition may dissuade best-practices units from sharing their best ideas (Butt et al. [2018], Li and Sandino [2018]). Even if best-practices units do share their successful ideas, learning units may not find them applicable or be able to replicate them. Furthermore, the

[3] Research shows that altruism (the perception of being helpful to others) and recognition (being recognized by and receiving feedback from others) are two of the main motivations for individuals to share information in social networks (Constant et al. [1994], Wasko and Faraj [2005], Brzozowski et al. [2009]).

knowledge creation literature states that a diverse range of ideas is an important source of knowledge development (Nonaka and von Krogh [2009]). A best practices initiative could backfire if it fixates the learning units' attention on a narrow set of ideas rather than on the wide range typically discussed on an ESN.

Based on these expected benefits and potential costs of promoting best practices on an ESN, we state our first hypothesis in null form:

> *Hypothesis 1: Introducing a best practices initiative on an ESN will not affect business unit financial performance.*

Given the potentially contingent nature of the effects of a best practices initiative, we examine four conditions that could moderate them: (a) the learning units' prior exposure to knowledge about best practices, (b) similarity in the markets served by the best-practices and learning units, (c) the learning units' prior performance, and (d) the learning units' prior trust in their manager's competence.

2.2. Potential Moderators of the Effectiveness of a Best Practices Initiative on an ESN

2.2.1. Prior Exposure to Knowledge

Research suggests that access to best practices via an ESN may lead to better results in units with limited physical exposure to other units' practices. For example, before having access to best practices via an ESN, store employees of a retail chain could learn ideas by visiting other nearby same-company stores and/or by engaging in discussions with the employees of those stores. However, employees in stores with fewer nearby same-company stores may have fewer opportunities for in-person visits and discussions, and hence less natural exposure to in-person best practices absent an initiative to share such practices via the ESN. Consistent with the notion

that stores with less natural exposure to ideas from other units would benefit more from an information sharing system, Li and Sandino [2018] found that introducing an information system through which retail employees shared marketing materials resulted in greater creativity improvements for the subsample of stores with fewer nearby same-company stores.

A follow up question is whether the mere presence of an ESN is enough to equalize employees' exposure to best practices. Prior studies suggest that this is not the case. For individuals already using an ESN, research shows that (a) the intensity of social connections in a social network significantly declines with geographic distance (Bailey et al. 2018), and (b) less-connected individuals are less likely to acquire knowledge from an ESN (Singh, Hansen and Podolny 2010). Singh, Hansen, and Podolny [2010] found that a reason for this is that less-connected employees are less able to seek help within the network. If that is the case, in the absence of a best practices initiative, employees from units with limited natural exposure to other units could also have limited exposure to useful information on the ESN. A best practices initiative could increase the visibility of useful information and help these employees understand which practices presented on an ESN would be most relevant and helpful to them. We therefore hypothesize:

> *Hypothesis 2: Introducing a best practices initiative on an ESN will lead to better financial performance in units with less exposure to same-company units than in units with more exposure to same-company units.*

Nevertheless, it is still possible that stores with more (rather than less) ex-ante exposure to same-company units would benefit more from a best-practices initiative. This could happen if discussions with employees from nearby units could help employees adapt ideas shared through the ESN best-practices initiative to their own context.

2.2.2. Similarity Between the Markets served by the Best Practices and Learning Units

Multiunit organizations often use ESNs to connect units serving different markets, thus exposing them to "newer/fresher" ideas. Along these lines, Li and Sandino [2018] found that introducing an information system to share creative work at a mobile phone retailer resulted in greater benefits for units operating in more divergent markets that required more customized and differentiated service. Exposed to a broader set of ideas, those units experienced increased creativity and employee engagement, consistent with the notion in organizational knowledge creation theory that more diverse ideas and experiences lead to greater innovation (Nonaka and Krough [2009]).

A best practices initiative on an ESN, however, is likely to direct the attention of learning units to the best-practices units' ideas, potentially resulting in exposure to fewer ideas overall than had they searched more broadly on the ESN. The ESN best practices initiative will likely have a more positive effect for learning units serving markets that are more similar to the markets served by the best-practices units, as a larger percentage of the knowledge shared will apply to them and the ideas shared may be easier to implement in their own markets. Indeed, Hansen [2002] and Levine and Prietula [2012] highlight that knowledge sharing only pays off if the practices shared by one unit can be reasonably useful to the units consuming that knowledge. We therefore hypothesize:

> *Hypothesis 3: Introducing a best practices initiative on an ESN will lead to better financial performance in units serving markets that are more similar to the markets served by the best practices units, than in units serving markets that are less similar to those markets.*

Despite the arguments stated above, the best practices initiative may still benefit units serving markets that are different from those served by the best practices units if it directs their attention to useful ideas they would not otherwise have considered.

2.2.3. Prior Performance

The initial performance of an organization is likely to condition the performance effects of introducing a best practices initiative on an ESN. On the one hand, lower-performing units might benefit the most from such an initiative, having more room for improvement and possibly more motivation to improve. On the other hand, lower performance may proxy for less capable or motivated employees or for worse market conditions, which may limit the ability to learn from a best practices initiative. For example, in a study examining the effects of online training (another form of online knowledge sharing), Fisher, Gallino, and Netessine [2018] find that employees in the second to highest performance quartile (above-average performers) are more likely to leverage new online knowledge than other employees.

In light of the different ways prior performance may moderate the effects of a best practices initiative on an ESN, we state our fourth hypothesis in null form:

> *Hypothesis 4: Introducing a best practices initiative on an ESN will not differentially affect financial performance in units with better prior performance than in units with worse prior performance.*

2.2.4. Trust in Management Competence

A best practices initiative may work more effectively if employees expect their managers to be more competent than other managers (what Das and Teng [2001] refers to as "competence trust"). That trust could encourage employees to seek support and guidance from their managers on how to translate best practices shared on the ESN into practice and results. However, it is possible that a best practices initiative would instead be less beneficial in such cases, since employees may rely more heavily on the ideas shared by their managers than on the best-practices ideas shared by their peers on the ESN. Furthermore, employees who trust their managers' competence may have

already sought advice on best practices from those managers, leaving less for the initiative to accomplish. Along those lines, Levine and Prietula [2012] find that the benefits of knowledge exchanges at a global consulting company were lower where the organization had alternative mechanisms to support learning. In the site they studied, alternative mechanisms supporting learning (such as broader access to on-site libraries, archives of past work, and online training), substituted for (rather than complemented) knowledge exchanges among peers. Whether employees' trust in management's competence increases or decreases the benefits of implementing a best practices initiative is, thus, an empirical question. We therefore state the following hypothesis in null form:

> *Hypothesis 5: Introducing a best practices initiative on an ESN will not differentially affect financial performance in units in which managers are trusted to be more competent than in units in which managers are perceived to be less competent.*

III. RESEARCH SETTING, SAMPLE, AND METHOD

3.1. Research Setting

Our study was conducted in collaboration with a large grocery store retailer operating a franchise model in multiple European countries. This retailer strongly empowered its franchisees, but also implemented structures that guided and limited the franchisees' actions (for example, it required the franchisees to keep a consistent store layout locating broad product categories in the same areas, to choose merchandise from a list of approved products, and to employ digital systems developed by the head office for personnel and inventory management). The retailer also shared best practices (from the top down) with its franchisees through a document that delineated

operating guidelines and a set of training sessions offered to existing and aspiring franchisees. The company had implemented an ESN similar to Facebook in one of the countries where it operated and had strongly encouraged the franchisees and staff in that country to use the ESN as their main method of communication. Employees on the ESN could form groups, create and share posts or photo/video albums, comment on posts, "like" or otherwise react to posts (e.g. with a "heart" or a "happy face" symbol), and send private messages to individuals within the network. Our study took place in this country.

The ESN was implemented a few years before this study began. It gave each user equal rights and privileges. A human-resources (hereafter, HR) manager mentioned in a conversation with the researchers that this was the first time when company personnel communicated directly to store employees, given that the employees were the franchisees' (rather than the company's) employees. At the beginning of 2019, the employees' engagement with the ESN was high, with more than 85 percent active. However, according to the HR managers, some users complained that the ESN was overloaded with information, some of it irrelevant to their work. Many ideas were shared on the network, but most were unremarkable. Furthermore, some HR and regional managers commented that many high-performing managers had expressed a preference for keeping their best ideas to themselves. To remedy these problems, the company issued rules and norms for interacting on the platform, shaping the ESN to be a tool used for constructive work purposes. But even when the company succeeded at promoting knowledge sharing to an extent (see the pre-intervention survey in Section 3.4 below), as in many companies implementing such networks, management was still concerned that the ESN could be distracting or overwhelming some employees rather than helping them share knowledge and improve job performance. These concerns were echoed in the pre-intervention survey, as described in Section 3.4 below.

To turn online engagement to offline performance improvements, the company decided to pilot-test a best practices initiative (hereafter, BPI) on its ESN in 2019 in collaboration with its regional management offices. The company had 12 regions with over 600 stores in the country where we conducted our study. The regional managers worked for the corporate headquarters, but each of the stores was managed independently by a franchisee who owned the store and kept the residual profit of the store s/he managed after paying a royalty to the company (franchisees were allowed to own only one store each).[4]

Before the study took place, one of the twelve regional offices operating in the country that we studied had already tried out an initiative promoting best practices on the ESN among its stores. This initiative was now modified and refined and then pilot-tested in five other regions (while the remaining six regions continued their usual use of the ESN without the BPI). The goal was to improve knowledge sharing and store operations to boost sales in the participating regions.

The company's HR team (in charge of managing the ESN) partnered with our research team to design and set up this BPI as a natural field experiment, which provided some significant advantages for research purposes. We were able to draw causal inferences, thanks to the random selection of the treatment stores, and to examine the effects of the BPI in a natural context (Bandiera et al. [2011], Floyd and List [2016]). Furthermore, because our study was a natural field experiment, subjects were unaware that they were participating in a study. Thus, we avoided self-selection and could discard alternative explanations related to the "Hawthorne effect".

The company implemented the BPI at its stores across these regions, conducted surveys to gather further information, and collected data on each store's normal course of operations. The company

[4] The company had a policy of granting only one store per franchisee (avoiding multi-unit franchising), because it wanted to encourage all franchisee-store managers to place their entire attention on their respective stores. The company enabled its best franchisees to switch to larger, more profitable stores when new stores became available, further increasing the managers' earnings potential.

later shared all of these data with our research team to facilitate our analysis of the effects of the BPI.

3.2. Sampling Strategy

3.2.1. Randomization Strategy

The treatment regions were selected with a stratified randomization strategy. The 11 regions used to test the BPI (as explained earlier, one of the 12 regions in the country had already experimented with a similar best practices initiative and was therefore excluded from the study) were split into three strata based on the store-level weekly sales over the past year: low sales trends (2 regions: 1 treated, 1 control), medium sales trends (5 regions: 2 treated, 3 control), and high sales trends (4 regions: 2 treated, 2 control). The regions were randomly assigned to the treatment and the control groups, stratum by stratum. The randomization strategy required us to bundle the two regions with the highest risk of "contamination" (Regions 2 and 8) into the same (treatment) group, as employees from these two regions regularly communicated with each other. The allocation of regions and stores into the treatment and control groups is shown in Table 1.

We collected data on store characteristics and the characteristics of the municipality in which the store was located (for example, demographics of the customer base in that area) and compared these characteristics between the treatment and the control groups. In Table 2, we show the balance of these characteristics between the two groups. Store age is the only characteristic that is statistically significantly different between the treatment and control groups: treatment stores tend to be younger. This difference is controlled for in our regression analyses estimating the treatment effects.

3.2.2. Sample Selection

The BPI was launched in late August 2019 and lasted until the end of December 2019. In our analyses, we define August 1, 2018 to the start of the intervention in August 2019 as the pre-intervention period (Post=0) and the last week of August to the end of January as the intervention period (Post=1). We selected a 12-month pre-intervention period to have enough data to estimate each store's individual sales trend while taking into account the full cycle of seasonal effects.

The company decided to focus on five product categories during the BPI, aiming to uncover the potential benefits of the BPI—by focusing on categories that accounted for a large majority of sales and for which the level of sales depended on the store teams' efforts—while minimizing the costs of implementation. As we will describe later, information about best practices was collected and organized by the regional office for each selected category; thus, the more categories covered, the greater the effort and cost. The five product categories were fresh goods (23.3% of sales), dry goods (19.5% of sales), beverages (11.9% of sales), fruits and vegetables (11.5% of sales), and bread (4.7% of sales)[5]. These categories account for a significant portion of the overall sales (the first four being the top four sales categories), have reasonable variation in sales across stores, and were considered by management most likely to benefit from the BPI.

We then conducted power analyses with simulated sales data (drawn from historical sales data); the results are shown in Appendix 1. These analyses aim to verify that the research design and planned tests would identify any meaningful effect of the BPI on the outcome of interest (sales). Our power analyses assumed a significance level of 10 percent (α=0.1; two-sided tests). They

[5] Fresh goods consist of goods that are distributed, stored and placed in a chilled environment (e.g. fish, meat, eggs, cheese). Dry goods mainly involve goods that do not require any special treatment in store (e.g. rice, canned food, chocolate). Beverages are typically liquefied drinking products (e.g. beer, soda, juice). Fruits and vegetables do not include canned fruits and vegetables. Bread are goods that are mostly daily fresh goods from the bakery, but also include some goods that are distributed frozen and sold in the Bread department.

show that we could be 80 percent confident that we would be able to identify effects of the BPI equal to or greater than a 1.3-percent change in total sales. The managerial team considered these minimal effect sizes to be "reasonable," given the effects of past sales initiatives. In our power analyses and subsequent formal analyses, we removed weeks that, according to the company, had historically been associated with extremely volatile sales (such as holidays and summer vacations), as these would add significant noise to the estimation and reduce the ex-ante power of the tests; the excluded weeks are listed in Appendix 1.

3.3. Details of the Intervention

The BPI began with a meeting of the general managers and sales managers from the regional offices of the five treated regions. The HR team managing the ESN explained the logic of the intervention and took questions and suggestions to ensure a smooth implementation. All of the managers involved welcomed the BPI enthusiastically. From August 26 to December 31, 2019, the managements of the five treatment regions featured and pinned best practices posts from their stores on the ESN. According to HR managers, these posts remained on the albums of the ESN in those regions after December 31, 2020. None of this was done in the six control regions. The timeline of the intervention is shown in Figure 1.

The HR team and the managers in each treatment region collaborated to select eight high-performing stores in August to be featured from August to November 2019. The 40 high-performing stores were selected based on proprietary information (including profitability measures

and other soft information) that were inaccessible to us.[6] Every month, two of each region's eight selected stores would be featured in a post. The two posts would appear two weeks apart. Each best practice post was pinned to the top of its corresponding regional group on the ESN for the two weeks during which it was featured, until the next best practice post replaced it. All best practice posts were kept in the group's archive (and were accessible to all store teams) once they were unpinned.

Regional offices were responsible for collecting, organizing, and posting the information about best practices. Each month, sales managers from the regional office visited the selected best practices stores (hereafter, BP stores), one store at a time, roughly two weeks apart. They interviewed the store managers about their best practices in the five featured product categories and did a "walk-through" of the store to take photos and videos of these best practices. Afterwards, the regional sales managers composed the best practices post for the regional group with texts, photos, and/or videos. These posts were in the form of an album that had the same format across all of the treatment regions, featuring a distinctive logo up front[7] (see Figure 2). In the main post, employees from the region could easily identify that this was a post about best practices from a selected high-performing store in the region. They could also see information about the store and its owner, a link to the album, and an overview of its photos and videos. Once they clicked on a photo, they could see a visual representation of the best practice (for example, how best to display

[6] Untabulated analyses show that the likelihood of a store being selected in the treated regions was significantly correlated ($p<=0.01$) with each of the performance measures that we use in our study (log sales, sales per employee, and sales per square meter), as measured in the pre-intervention period. Regarding the implementation of the BPI, 36 of the 40 selected stores followed the intervention schedule to make best practices posts, but four stores (from three regions) did not make the best practices posts by the end of the intervention period (December 31, 2019). The reasons provided were that they had conflicts with preparations for the holiday season and (in two cases) that, according to the store managers, the main best practices they planned to share had already been covered by other stores who posted earlier.

[7] In the figure, we replaced the logo of the company with a green generic logo to protect confidentiality.

bananas), the title of the best practice, a description of its benefits, and an explanation of it (what it means, how it is done, additional tips, and so on). These posts were visual and easy to understand.

Appendix 2 presents a few representative posts from the BPI. The vast majority of posts presented an appealing picture of a product/aisle display with an explanation of the decisions that the store staff made for those products or displays and the results obtained. In some occasions, management emphasized the relevance of empowering employees to make sure they took ownership of product decisions, for example, understanding customer preferences that should be considered when choosing what products to offer, taking care of product-related challenges (e.g., those arising from mixing up goods with different expiration dates), and searching for opportunities to enhance product handling to serve the best interests of the store. In a handful of cases we observed potentially conflicting posts. For instance, while the fruits-and-vegetables post in Appendix 2 (post 5, from a supermarket in an urban area) recommended positioning best-selling products in all the counters, and in a way that would "bring up the volume and growth of bestsellers," a fruits-and-vegetables post from a different supermarket located in a less-densely populated area, recommended placing best-seller products on "less" attractive spaces (the logic offered by this store was that customers would search for those items regardless of where they were placed). This suggested that stores had different perspectives on what constituted best practices, potentially reflecting local customer preferences. The posts received a fair amount of attention. Within each treated region, we tracked the "likes" and comments for each best-practices post and the number of users who had seen them. On average, within two weeks of posting, a BPI post received 33.5 "likes" (up to as many as 100) and close to three comments (up to as many as 22) and was seen by 489 users (up to as many as 830).

There were some advantages in having the regional office collect best practices information and create the posts. First, it motivated the selected high-performance stores to share their best practices. Second, it led to the adoption of a consistent protocol across all the treatment regions for selecting high-performing stores, gathering best practices information, and creating the best practices posts. Although this limited the format of the BPI and imposed some cost on the regional office, it improved the quality of the treatment in this natural field experiment.

3.4. Pre- and Post-intervention Surveys

We designed and asked the company to implement a pre- and a post-intervention survey to help us collect more data about the employees' experience at work and using the ESN. A key element that we wanted to assess and incorporate in our analyses (specifically to test H5, as explained later) was the trust store employees had on the regional manager's competence.

Our purpose in asking about the employees' experience of using the ESN was to clarify why and how they used the ESN, and whether they were finding relevant knowledge in it. In total, employees from 491 stores in the 11 regions analyzed provided responses to the pre-intervention survey.[8] 85% of those store responses indicated that they had used the ESN system for more than 5 minutes/week over the past two months. We asked those employees for their top three reasons for *following* content on the ESN. The main reasons were (1) learning how to do a better job (34% of responses); (2) learning how to have a better and more fulfilling work-life (23% of responses); and (3) feeling like part of the organization (21% of responses).[9] We also asked whether they had

[8] These 491 responses were entered by one or two individuals per store. We averaged out responses where more than one individual entered a response for a store.

[9] *The following question was asked:*
I follow information on ESN (e.g., posts, photos, comments) because… (please select up to three reasons)
__I want to know others who work at the Company.
__I want to learn useful information that would help me do a better job.
__I am curious to learn more about my colleagues.

contributed content on the ESN system, and, if they did, why. The motivations most frequently given for *sharing* content on the ESN were: (1) helping other employees do a better job (28% of responses); (2) contributing one's fair share as other employees had done (22% of responses); and (3) seeking feedback or solutions (20% of responses).[10] It was unclear whether a BPI would suppress some employees' motivations to share content as the BP posts could reduce their need to help others. Others could have perceived that they would not need to contribute a fair share but would need to wait for the regional office to select them for a BP post. Yet others could have taken the BPI as a signal that only posts from high performing stores would be valued.

The pre-intervention survey suggested that employees appreciated the information that they gathered from the ESN. Using a Likert scale where 1 = never, 2 = rarely, 3 = sometimes, 4 = very often, and 5 = always, on average, employees indicated that they agreed "very often" with the statements "It is easy to find the information I need on the ESN" (mean=3.6, median=4) and "The information on Workplace helps me do my job better" (mean=3.8, median=4). Employees applying ideas that they saw on the ESN indicated that the top three characteristics of the ideas that they adopted were that they were (1) creative or novel (36% of responses); (2) posted by individuals that had already generated positive results (29% of responses); and (3) easy to copy (13% of responses). [11]

__I want to learn information that can help my career prospects at the Company.
__I want to learn information that could help me live a better and more fulfilling work-life at the Company.
__I want to feel a part of the organization I work for.

[10] *This question was stated as follows:*
I share content on the ESN because . . . (please select up to three reasons)
__my ideas could help other employees at the Company to do a better job.
__I want the Company employees to know what I am good at.
__other Company employees have contributed useful content to the system and I want to contribute my fair share as well.
__getting recognition from colleagues motivates me.
__I want to ask for feedback or solutions from other employees.
__most other employees post content and I feel the pressure to do so as well.
__the regional office encourages sharing content on Workplace.

[11] *The following question was asked:*

Despite the evidence above suggesting that employees used the ESN to share knowledge, we also confirmed the concerns that led management to implement the BPI. Using the same Likert scale from 1 = never to 5 = always, respondents indicated, on average, that they agreed "sometimes" to "very often" with the statement "I am *overwhelmed* by the amount of information on [the ESN]" (mean=3.5, median=4); and "sometimes" with the statement "I cannot decide which information is useful for me on [the ESN]" (mean=2.9, median=3).

3.5. Descriptive Statistics and Visualization of Treatment Effects

Our final dataset has 32,564 store-week observations (containing the sales of the five product categories). Table 3 shows the descriptive statistics of our sample data. Our main dependent variable is the natural log of sales. Sales, reported in US dollars (rather than the local currency, to protect the confidentiality of the company), vary widely across store-weeks. An average store sells US$97,626 in groceries each week, though weekly store sales range from US$ 0.04 to US$ 319,316 across store-weeks.[12] The observations from the treatment group account for 46 percent of the store-week observations. We use *Average Number of Reactions* and *Average Times Seen* to measure the popularity of the BP posts. Specifically, these two measures capture the number of reactions ("likes", "hearts", etc.) to a BP post and the number of users who have seen a BP post within 2 weeks of the initial post, averaged across all of the BP posts that a focal store was exposed to over the post-intervention period. *Average Number of Reactions* to BP posts ranges from 22.5

The ideas from the ESN system during the last two months that I have applied are because . . .
(please select up to three reasons)
__they are creative or novel.
__they are easy to copy.
__they have been shared by top-performing stores.
__they have been shared by people I trust.
__they have many "likes."
__the ones who shared the idea have shown they have already generated positive results.

[12] Our main results in Table 5 are robust to dropping six outlier store-weeks with sales below US$10.

to 65 (with a mean value of 35.5), while *Average Times Seen* for BP posts ranges from 318 to 743 (with a mean value of 513), among the treatment stores.

To test Hypotheses 2, 3, 4, and 5, we constructed four moderating variables: *# of Nearby Stores*, *Proximity to BP Stores*, *Prior Performance*, and *Trust in Regional Manager Competence*. Later on, we turn each moderator into (a) a dummy variable indicating whether or not the underlying raw measure was above the sample median, and (b) a normalized z-score version of the measure (where we subtracted the mean from each raw measure, and divided it by its standard deviation). We do not include these transformed measures in Table 2 but use them in our regression analyses for ease of interpretation (i.e., Tables 5-9).

# of Nearby Stores is the number of same-company stores within a 10-kilometer radius of the focal store. The average number of nearby stores for the stores in our sample is 16, ranging from 0 to 78 nearby stores. *Proximity to BP Stores* is equal to 200 minus the average # of kilometers between the focal store and the BP stores in the region where the focal store was located (we picked 200 kilometers since the longest average distance between any focal store and its regions' BP stores was just under this distance). The average distance between the treatment stores and the BP stores is about 38 kilometers, hence, our measure for Proximity to BP stores has a mean value of 162 (200-38). *Prior Performance* is a store's average value of weekly sales per employee over the pre-intervention period. We find that the average weekly sales per employee amount to US$ 5,659, varying widely from US$ 1,962 to US$ 65,747 across all store-weeks. *Trust in Regional Manager Competence* is the principal component from three pre-intervention survey questions related to the regional manager. Respondents were asked to indicate, on a Likert scale from 1 (never) to 5 (always) (including a "Not Applicable" option if they did not know the regional manager), the extent to which each of the following statements described their experience at work: (1) "The

regional manager is very capable of performing her/his job;" (2) "The regional manager is known to be successful at the things he/she wants to implement;" and (3) "I feel confident about the regional manager's skills." The principal component explains 88.8 percent of the variation in the three questions. The loadings of Questions 1–3 into the principal component measure are 0.580, 0.572, and 0.581, respectively. This measure also varies widely, from -1.19 to 0.92.

Before testing the effects of the BPI on financial performance, we examine whether the BP posts successfully captured employees' attention as expected. In Table 4, we present difference-in-differences regressions suggesting that the initiative indeed led to (1) more posts and comments (made by store employees) in the regional groups where the BPI was implemented (as shown by the results in Columns 1 and 3); (2) more posts and comments (made by store employees) in the regional groups relative to the total number of posts and comments made in both the store and the regional groups (as shown by the results in Columns 2 and 4). This increase in posts and comments made by store employees (beyond the BP posts entered by the regional sales representative) suggests that rather than discouraging employees from sharing additional content, the BPI triggered greater engagement in *inter-store* knowledge sharing leading to additional posts and comments.

Next, we visualize the treatment effect by plotting the adjusted natural log of weekly sales (*Ln(Sales)*) against time (in weeks), contrasting the treatment and control groups. In Figure 3, the vertical axis is the residual of regressing *Ln(Sales)* on store fixed effects, week fixed effects, and store-time-trend fixed effects.[13] The horizontal axis represents time (in weeks). Prior to the BPI, treatment and control groups exhibited similar trends in sales. Following the BPI, the sales trends

[13] Similar to Deller and Sandino (2020), in this study we assess the effect of the system intervention (BPI) on the natural logarithm of sales because (a) we aim to analyze the effect of the BPI on the percentage change in sales, given that the initial level of sales differed across stores, and (b) the natural logarithm of sales was closer to being normally distributed, in line with the necessary assumptions required to run OLS regressions.

for treatment and control groups gradually diverge, with the treatment group showing a relatively favorable sales trend over time. This is also consistent with our expectations: it would take some time for such best-practices posts (and related discussions) on the ESN to improve offline performance outcomes such as sales.

3.6. Research Design

We examine changes in the sales performance of the treated stores relative to that of the control stores. Motivated by the visualization of the treatment effects in Figure 3 and our desire to capture the sales-trend effect associated with the BPI, rather than focusing on an average treatment effect on sales *level*, we add the time dimension and focus on the effect on sales *trend*.

We estimate the following model:

$$Ln(Sales)_{it} = \delta_0 + \delta_1 Post_t + \boldsymbol{\delta_2\, Best\ Practices_i \times Post_t} + \delta_3 Time_t$$
$$+ \delta_4 Time_t \times Post_t + \boldsymbol{\delta_5\, Best\ Practices_i \times Post_t \times Time_t} + \delta_L \text{ Store Fixed Effects}$$
$$+ \delta_M \text{ Week Fixed Effects} + \delta_N \text{ Store Trend Fixed Effects} + \varepsilon_{it},$$

where $Ln(Sales)_{it}$ is the natural log of sales for store i on week t; $Post_t = 1$ if week t is or comes after the first week of the BPI (August 26); $Time_t$ is the number of weeks relative to the first week of the initiative (-52 to +22, with 0 being the first week);[14] and $Best\ Practices_i = 1$ if store i is a treatment store. The main focus of our estimation is the *sales-trend effect* (δ_5); that is, the change in sales relative to that of control stores for each passing week since the beginning of the BPI. Therefore, $\delta_2 + Time_t * \delta_5$ will be the total *Sales Effect as of Time t*; that is, the change in sales relative to that of control stores by the end of week t. Standard errors are clustered by store. Notice that the 36 BP stores are also likely to benefit from the intervention in the periods when they are

[14] The last round of BP posts was put into the ESN in December 2019 (18 weeks after the start of the initiative). We included four additional weeks of sales data (until the end of January 2020) in our analyses because it would take time for people to learn from and apply the ideas from these BP posts.

not the source of ideas. Yet, we conduct robustness tests that replicate all of our analyses excluding these 36 stores from the sample.

IV. ANALYSIS AND RESULTS

4.1. Effects of the ESN Best Practices Initiative on Financial Performance

First, we test Hypothesis 1 by estimating the above model in our full sample and examine the effect of the BPI on sales and on the sales trend. In Table 5, the interaction term between *BP* and *Post* has a positive but statistically insignificant coefficient (0.0034), indicating that the intervention had little effect on sales at the beginning of the intervention (week 0). However, δ_5 (the coefficient for $BP \times Post \times Time$) is positive and significant (at the 5% level), indicating that the BPI had a positive and significant effect on sales with each passing week. If we multiply this coefficient by 18 and take the exponential, we can see that the sales-trend effect of the 18-week intervention resulted in sales of the treatment group being higher than those of the control group by 3.67 percent ($e^{(0.002*18)} - 1$).[15] Although we were not able to predict definitively prior to the intervention whether the BPI would have any effect on performance, our results show that it did (rejecting our null hypothesis H1).[16] However, the effect took some time to manifest. This is consistent with the nature of a BPI on an ESN: it takes time for store managers and employees to read and understand the posts, discuss them, decide which best practices to implement, implement them, and wait to see if the idea actually works. Our results suggest that the BPI lead to an increase in the rate at

[15] The total effect, including the immediate effect and the learning effect of the intervention ($e^{(0.0034+0.002*18)} - 1$), indicates that the treatment group's sales where 4.02 percent higher relative to the control group by the end of the treatment period.

[16] This result, supporting H1, as well as all our subsequent results reported below (from H2 to H5) are robust to excluding the 36 stores which posted the best practices (BP stores).

which employees learned (i.e., identified and applied relevant ideas) from the content shared on the ESN.

To better understand whether the performance improvements are related to the extent to which employees accessed the BP posts, in Table 6, we examine whether the popularity of these BP posts (or the amount of attention paid to the posts) moderates the performance effects of the BPI. We use *Average Number of Reactions* and *Average Times Seen* to measure the popularity of the BP posts. Specifically, these two measures capture the number of reactions ("likes", "hearts", "happy faces," etc.) per BP post and the number of users who saw a BP post within 2 weeks of the initial post averaged across all of the BP posts in the region (that is, all of the BP posts that a focal store was exposed to over the treatment period).

In Table 6 Columns 1 and 2, we ran our baseline regressions in the two subsamples where stores were exposed to relatively less popular posts and more popular posts (as measured by either the average number of reactions to the BP posts or the average number of users who saw the posts – these two regional-level measures generate the same subsamples). We find that the subsample of stores exposed to more popular posts experienced a positive and significant sales trend effect (the coefficient on $BP \times Post \times Time$ is 0.0044 and significant at the 1% level; that is, learning from the BP posts resulted in 8.24 percent higher sales in these stores than the control group by the end of the 18-week initiative) while the subsample of stores exposed to less popular posts didn't experience any significant sales-trend effect relative to the control group. In Column 3, we combine the two subsamples and interact the binary indicator for being exposed to more popular posts with the sales-trend effect term ($BP \times Post \times Time$) and find that the difference in the sales-

trend effect between the two subsamples is statistically significant.[17] In Columns 4 and 5, we use the underlying raw measures for post popularity, *Average Number of Reactions* and *Average Times Seen*, as continuous-value moderators, and find positive and significant interaction coefficients between the moderator and the sales-trend effect term (*BP* × *Post* × *Time*), indicating that the sales-trend effect from the BPI became stronger as the amount of attention paid to the BP posts increased.[18]

4.2. Factors Moderating the Effectiveness of the Best Practices Intervention

Next, we delve deeper into the cross-section of stores and examine circumstances potentially moderating the effects of the intervention.

To test Hypothesis 2, we use *# of Nearby Stores* as the moderator. Prior to the BPI, stores with fewer nearby same-company stores were likely to have less best practices knowledge gained through natural exposure to ideas from other company stores. A BPI on the ESN, as a result, would be more useful to these stores. In Table 7, the first two columns show the regression results for the two subsamples resulting from splitting the stores based on the number of nearby same-company stores (measured by number of stores within a 10-kilometer radius of the focal store) in their region. In the subsample with fewer nearby stores, the BPI had a positive and statistically significant effect on the sales trend (0.0029; that is, after 18 weeks, the BPI was associated with a sales-trend effect explaining 5.36-percent higher sales in the treatment group relative to the control

[17] Notice that in this particular table, the number of observations in the subsamples in the first two columns do not add up to the total observations in the three right-most columns. This is because the moderator measures (*Popular Posts_High, Average Number of Reactions,* and *Average Times Seen*), having been constructed in relation to the BP stores, could only be estimated for the treatment stores. We therefore included *all* observations corresponding to the control stores in the subsamples in the first two columns.

[18] We replicated the analyses in Table 6 using the number of comments made to the BP posts in the first 2 weeks after the posts were entered in the ESN. We do not tabulate these analyses given that the average number of comments per post was three. Nevertheless, our results are consistent with those reported in Table 6: our sales-trend effect is driven by posts receiving a higher number of comments.

group). In contrast, the BPI did not seem to have an impact on sales in the subsample with more nearby stores. These results are consistent with H2. In Columns 3 and 4, we combined the two subsamples and examined the coefficient on the four-way interaction term (*BP* × *Post* × *Moderator* × *Time*), where *Moderator* was either (a) a dummy variable indicating that the # *Nearby Stores* was above median (Column 3), or (b) a z-score transformation of the # *Nearby Stores* variable (Column 4). The coefficient is negative in both cases, consistent with the earlier comparison between the two subsamples; that is, there was a less positive effect on the sales trend (less positive learning effect) for the stores with greater best-practice knowledge prior to the BPI.[19] However, this result was insignificant. The results in Table 7 are only partially consistent with our hypothesis (H2): learning units with less exposure to knowledge prior to the initiative (that is, with fewer nearby same-company stores) experienced greater sales improvement—but not significantly greater—than those with greater prior exposure to knowledge.

To test Hypothesis 3, we use *Proximity to BP Stores* to measure the degree to which the markets served by the BP stores resemble those served by the focal store. In Table 8, the first two columns show the regression results in the two subsamples that showed relatively less and relatively greater proximity to the BP stores in their region.[20] The intervention had positive and significant effects on the sales trend only in the subsample where the stores were closer to the BP stores. The magnitude of the coefficient (0.0034) for this subsample suggests that the BPI sales-trend effect resulted in 6.31-percent higher sales in the treatment group relative to the control group by the end of the 18-week initiative). In Columns 3 and 4, we combine the subsamples and examine the

[19] Results are similar when we define "nearby stores" as stores within a 5-kilometer or 3-kilometer radius of the focal store.

[20] As in Table 6, in Table 8 the number of observations in the subsamples in the first two columns do not add up to the total observations in the two right-most columns the moderator measures (Proximity to BP Stores, Proximity to BP Stores_High) could only be constructed in relation to the BP stores.

coefficient on the four-way interaction term (*BP* × *Post* × *Moderator* × *Time*), where *Moderator* was either (a) a dummy variable indicating that the # *Proximity to BP Stores* was above median (Column 3), or (b) a z-score transformation of the # *Proximity to BP Stores* variable (Column 4). The coefficient on the quadruple interaction proves to be negative and statistically significant in both cases, consistent with the earlier comparison between the two subsamples. In other words, the stores with greater proximity to the BP stores experienced less improvement in sales trend.

The results in Columns 3 and 4 of Table 8 are consistent with Hypothesis 3. We find that the performance effect of the BPI differs depending on the similarity between the markets served by the learning units and those served by the BP units. Specifically, learning units serving markets in closer proximity to those served by the BP units experienced greater sales growth.[21] Despite stores serving different markets than the BP units could have found the best practices to be more novel, a BPI is more beneficial to units serving markets in closer proximity to those served by the BP units. This suggests that the applicability of the shared knowledge is key to driving the learning units' financial performance. We note that a focal store located very close to the BP stores could have learned those BP stores' practices prior to the BPI, resulting in less potential to benefit from the initiative. In untabulated analyses, we reran our analyses in Columns 3 and 4 of Table 8 excluding any stores that were within a 10 km radius of the BP stores. Consistent with our intuition,

[21] As a robustness check, we use an alternative measure for the similarity between the markets served by the focal store and the BP stores. This alternative measure is *Market Divergence* (an opposite measure of market similarity), which is an aggregate measure of the difference between the demographic characteristics in the municipality of the focal store and the average demographic characteristics of the municipalities of the BP stores in the same region. We consider five dimensions (population density, average age of the population, household size, per-capita income, and education level) and use the methodology used in Campbell, Datar, and Sandino 2009 to estimate the divergence measure as a sum of the normalized differences in the values of these dimensions between the focal store and the BP stores. The results are similar; the BPI was associated with greater financial performance when the markets served by the learning units and the BP units were more similar to each other, i.e. exhibiting lower *Market Divergence*. However, the results for the quadruple interaction (equivalent to those in the third and fourth columns of Table 8) were significant for the dichotomous, but not for the continuous variable.

the exclusion of those stores resulted in an even stronger effect of the *Proximity to BP stores* on the magnitude of the sales-trend effect associated with the BPI.

To test Hypothesis 4, we use prior performance as the moderator. In Table 9, the first two columns show the regression results in the two subsamples including stores with lower versus higher prior performance. We measure prior performance as the average sales per employee for the store-weeks prior to the BPI. In the subsample with lower-than-median prior performance, the BPI had a positive and statistically significant effect on the sales trend (0.0034; that is, after 18 weeks, the BPI sales-trend effect resulted in 6.31-percent higher sales in the treatment group relative to the sales of the stores in the control group). The BPI had an insignificant effect on sales when the stores had above-median prior performance. In Columns 3 and 4, we combine the two subsamples and examine the coefficient on the four-way interaction term (*BP* × *Post* × *Moderator* × *Time*), where *Moderator* was either (a) a dummy variable indicating that *Prior Performance* was above median (Column 3), or (b) a z-score transformation of the *Prior Performance* variable (Column 4). The coefficient on this interaction term is negative and statistically significant regardless of the moderator variable used, consistent with the earlier comparison between the two subsamples; that is, there was a greater sales-trend effect for stores with lower prior performance. The results in Table 9 reject our null hypothesis 4: we find that learning units with *worse* pre-intervention performance experienced greater sales growth, consistent with the idea that these units had more room for improvement.[22]

To test Hypotheses 5, we use *Trust in Regional Manager Competence* as the moderator. In Table 10, the first two columns show the regression results in the two subsamples with relatively lower versus higher employee trust in the regional manager competence prior to the BPI. In the lower-

[22] As a robustness check, we define prior performance as sales scaled by the store's physical area (in square meters). The results are qualitatively similar to those reported in Table 9.

trust subsample, the BPI had a positive and statistically significant effect on the sales trend (0.0043; that is, after 18 weeks, the BPI resulted in 8.05-percent higher sales in the treatment group than in the control group). In contrast, the BPI had an insignificant effect on sales in the high-trust subsample. In Columns 3 and 4, we combine the two subsamples and examine the coefficient on the four-way interaction term (*BP* × *Post* × *Moderator* × *Time*), where *Moderator* was either (a) a dummy variable indicating that *Trust in Regional Manager Competence* was above median (Column 3), or (b) a z-score transformation of the *Trust in Regional Manager Competence* variable (Column 4). The coefficient is negative in both cases (although only statistically significant in Column 3), consistent with the earlier comparison between the two subsamples; that is, the effect on sales trend was more positive for stores with lower trust in the regional manager prior to the BPI.

The results in Table 10 reject the null hypothesis 5: we find that learning units where employees had *lower* trust in the regional manager competence experienced greater improvements in sales. This somewhat counterintuitive result suggests that, for driving unit performance, a BPI on an ESN could serve as a substitute mechanism for trust in management competence. This is consistent with a prior study suggesting that knowledge transfers among co-workers are less beneficial in organizations already providing greater learning support (Levine and Prietula 2012).

V. CONCLUSION

We use data from a natural field experiment in a large retail chain to examine the effects of a BPI via an ESN. Our results do not show an immediate improvement in sales, but do show an improvement in sales trends in the treatment group relative to the control group, amounting to a 3.67-percent sales increase by the end of the four-month initiative. This effect was stronger when the units were exposed to posts that received relatively more attention. Furthermore, the effects of

the initiative were more positive in stores with *less* prior exposure to best practices knowledge (that is, with fewer nearby same-company stores), in stores that served markets that were more similar to the markets of the best-practices units (that is, stores to which the knowledge shared was more relevant), in stores that had had lower performance before the intervention, and in stores whose employees, before the BPI, trusted their regional manager's competence *less* rather than more. Overall, these results suggest that a best practices initiative on an ESN could improve sales, but that such results depend on a given unit's condition.

While our results are robust to alternative specifications, they should nevertheless be interpreted with caution. First, because the company implemented the initiative following the protocol of a natural field experiment, it confronted limitations that could have reduced the power of the tests. For example, to ensure treatment quality, we did not give individual regions or stores the freedom to customize the content, format, or timing. Second, the organization we studied had an engaging culture, potentially contributing to the effectiveness of its ESN. The BPI effect we document may not be generalizable to organizations with less engaging cultures. Future research can therefore examine whether the effect of a best practices initiative on an ESN would depend on complementary organizational practices or elements of organizational culture. Third, the long-term learning effects of a BPI (beyond the 18-week period that we analyzed) remain unclear. Future research can expand the duration of similar interventions to explore longer-term effects.

Although we examine how best practices initiatives on ESNs improve performance through better knowledge sharing, it is possible that performance improvements occur because learning units are motivated to improve performance to increase their own chances of being recognized as a best practices unit in the future and being featured on the ESN. Future research can explore this incentive channel in addition to the information-sharing mechanism.

REFERENCES

Bandiera, O., Barankay, I., & Rasul, I. (2011). Field experiments with firms. *Journal of Economic Perspectives*, 25 (3), 63–82. https://doi.org/10.1257/jep.25.3.63

Brzozowski, M. J., Sandholm, T., & Hogg, T. (2009). Effects of feedback and peer pressure on contributions to enterprise social media. *GROUP'09 - Proceedings of the 2009 ACM SIGCHI International Conference on Supporting Group Work*, *June 2014*, 61–70. https://doi.org/10.1145/1531674.1531684

Butt, M. N., Antia, K. D., Murtha, B. R., & Kashyap, V. (2018). Clustering, knowledge sharing, and intrabrand competition: A multiyear analysis of an evolving franchise system. *Journal of Marketing*, 82 (1), 74–92. https://doi.org/10.1509/jm.16.0173

Campbell, D., Datar, S., & Sandino, T. (2009). Organizational design and control across multiple markets: The case of franchising in the convenience store industry. *Accounting Review* 84 (6): 1749-1779.

Campbell, D., Epstein, M. J., & Asis Martinez-Jerez, F. (2011). The learning effects of monitoring. *Accounting Review*, 86 (6), 1909–1934. https://doi.org/10.2308/accr-10129

Campbell, D., Erkens, D. H., & Loumioti, M. (2014). Exception Reports as a Source of Idiosyncratic Information. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.2482901

Charki, M.-H., Boukef, N., & Harrison, S. (2018). *Maximizing the Impact of Enterprise Social Medi*. MIT Sloan Management Review. https://sloanreview.mit.edu/article/maximizing-the-impact-of-enterprise-social-media/

Cheng, M. M., & Coyte, R. (2014). The effects of incentive subjectivity and strategy communication on knowledge-sharing and extra-role behaviours. *Management Accounting Research*, 25(2), 119–130. https://doi.org/10.1016/j.mar.2013.07.003

Constant, D., Kiesler, S., & Sproull, L. (1994). What's mine is ours, or is it? A study of attitudes about information sharing. *Information Systems Research*, *5*(4), 400–421. https://doi.org/10.1287/isre.5.4.400

Costello, K. (2019). *Gartner says worldwide social software and collaboration revenue to nearly double by 2023*. https://www.gartner.com/en/newsroom/press-releases/09-24-2019-gartner-says-worldwide-social-software-and-collaboration-revenue-to-nearly-double-by-2023

Das, T. K. and B. Teng. (2001). Trust, Control, and Risk in Strategic Alliances: An Integrated Framework. *Organization Studies,* 22 (2): 251-283.

Deller, S. and T. Sandino (2020) Effects of a Tournament Incentive Plan Incorporating Managerial Discretion in a Geographically Dispersed Organization. *Management Science* 66(2), 911-931. https://doi.org/10.1287/mnsc.2018.3214

Fisher, M., S. Gallino, and S. Netessine. (2018). Does Online Training Work in Retail? *SSRN Electronic Journal*: http://dx.doi.org/10.2139/ssrn.2670618

Floyd, E., & List, J. A. (2016). Using Field Experiments in Accounting and Finance. *Journal of Accounting Research*, 54 (2), 437–475. https://doi.org/10.1111/1475-679X.12113

Hansen, M. T. (2002). Knowledge networks: Explaining effective knowledge sharing in multiunit companies. *Organization Science*, 13 (3), 232–248. https://doi.org/10.1287/orsc.13.3.232.2771

Hwang, Y., Erkens, D. H., & Evans, J. H. (2009). Knowledge sharing and incentive design in production environments: Theory and evidence. *Accounting Review*, *84*(4), 1145–1170. https://doi.org/10.2308/accr.2009.84.4.1145

Kuegler, M., Smolnik, S., & Kane, G. (2015). What's in IT for employees? Understanding the relationship between use and performance in enterprise social software. *Journal of Strategic Information Systems*, *24*(2), 90–112. https://doi.org/10.1016/j.jsis.2015.04.001

Leonardi, P. M. (2014). Social media, knowledge sharing, and innovation: Toward a theory of communication visibility. *Information Systems Research*, *25*(4), 796–816. https://doi.org/10.1287/isre.2014.0536

Leonardi, P. M. (2018). Social media and the development of shared cognition: The roles of network expansion, content integration, and triggered recalling. *Organization Science*, *29*(4), 547–568. https://doi.org/10.1287/orsc.2017.1200

Levine, S. S., & Prietula, M. J. (2012). How knowledge transfer impacts performance: A multilevel model of benefits and liabilities. *Organization Science*, *23*(6), 1748–1766. https://doi.org/10.1287/orsc.1110.0697

Li, S. X., & Sandino, T. (2018). Effects of an Information Sharing System on Employee Creativity, Engagement, and Performance. *Journal of Accounting Research*, *56*(2), 713–747. https://doi.org/10.1111/1475-679X.12202

Neeley, T. B., & Leonardi, P. M. (2018). Enacting knowledge strategy through social media: Passable trust and the paradox of nonwork interactions. *Strategic Management Journal*, *39*(3), 922–946. https://doi.org/10.1002/smj.2739

Nonaka, I., & von Krogh, G. (2009). Tacit knowledge and knowledge conversion: Controversy and advancement in organizational knowledge creation theory. *Organization Science*, *20*(3), 635–652. https://doi.org/10.1287/orsc.1080.0412

Oleson, J. (2013). *The Rise of the Enterprise Social Consultant*. Wired. https://www.wired.com/insights/2013/11/the-rise-of-the-enterprise-social-consultant/

Singh, J., Hansen, M. T., & Podolny, J. M. (2010). The world is not small for everyone: Inequity in searching for knowledge in organizations. *Management Science*, *56*(9), 1415–1438. https://doi.org/10.1287/mnsc.1100.1201

Wasko, M. M. L., & Faraj, S. (2005). Why should I share? Examining social capital and knowledge contribution in electronic networks of practice. *MIS Quarterly: Management Information Systems*, *29*(1), 35–57. https://doi.org/10.2307/25148667

Figure 1: Timeline

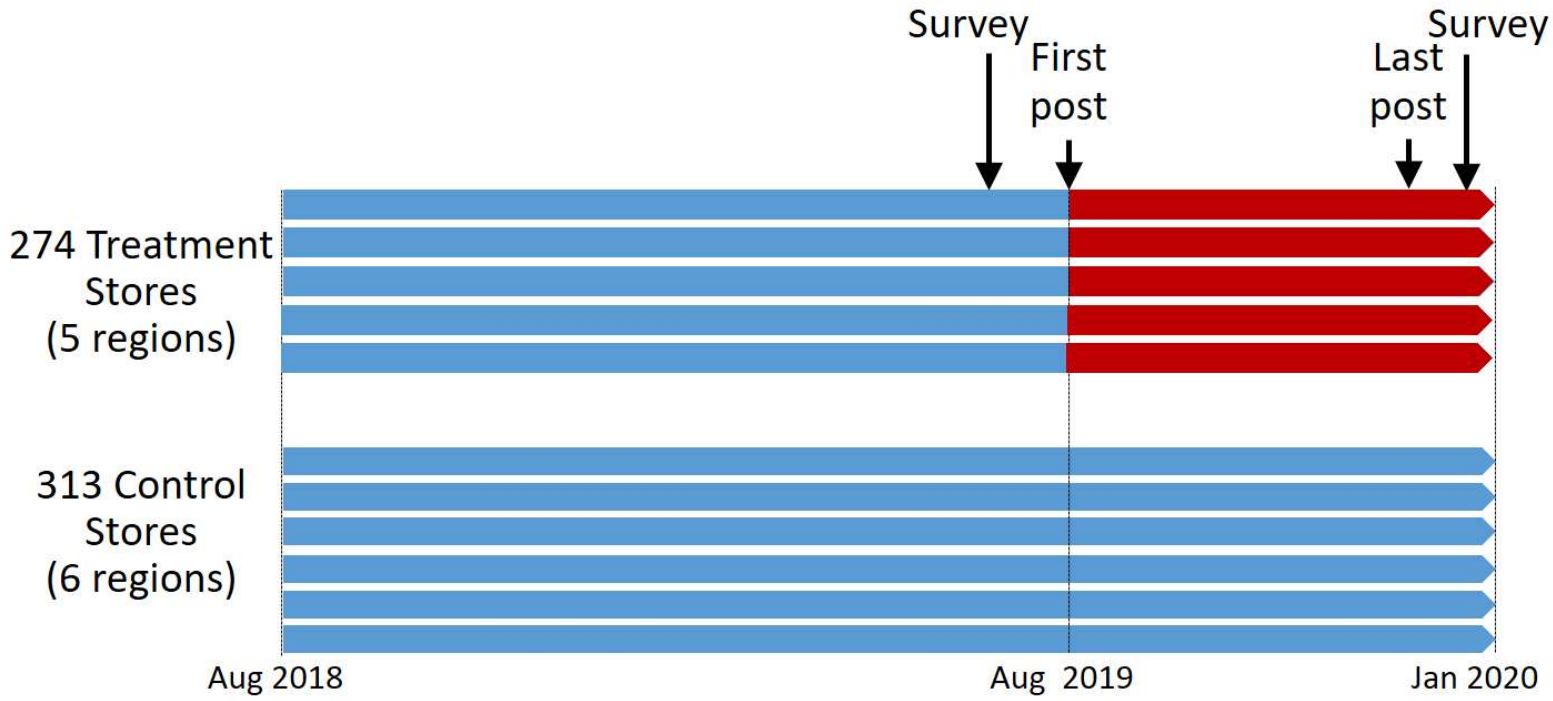

Figure 2: Best Practice Posts

Figure 3: Visualization of Treatment Effects

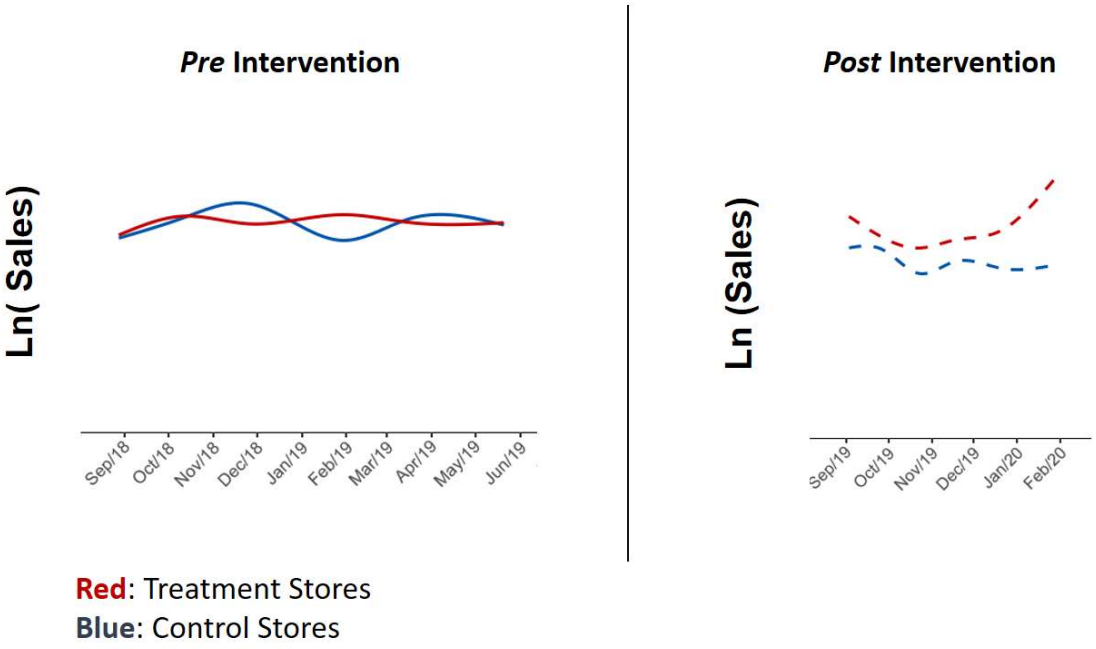

The vertical axis is the residual of regressing ln(Sales) on store fixed effects, week fixed effects, and store-time-trend fixed effects. The horizontal axis represents time (measured in weeks).

Table 1 Stratified Randomization Outcomes Allocating Treatment and Control Regional Groups within Three Sales-trend Strata

	Best practices (treatment) group		**Control group**	
		# stores		# stores
Low sales trends	Region 1	37	Region 2	53
Medium sales trends	Region 3	70	Region 5	51
	Region 4	55	Region 6	68
			Region 7	68
High sales trends	Region 8	57	Region 10	38
	Region 9	55	Region 11	35
	Total	274	*Total*	313

Table 2 Covariate Balance between Treatment and Control Groups[23]

	Best Practices (Treatment) Group	Control Group	t-test p-value
Varying by stores:	274 stores	313 stores	
Gross area (square meters)	1,296.56	1,309.71	0.69
Net area (square meters)	880.16	900.82	0.36
Store age in years	14.72	16.40	**0.05**
Hours open during the week	15.86	15.96	0.15
Weekly sales, 2018 (in US$)	137,007	142,400	0.25
Weekly sales trend, %	0.07	0.09	0.30
Population density (2018)	520.18	327.44	0.42
Average age (2018)	39.94	39.52	0.62
Average household size (2018)	2.15	2.22	0.36
Average household income (2017) (in US$)	55,357	57,802	0.12
% with secondary(+) education (2017)	0.74	0.74	0.97
Average store counts, by muni	20.48	11.13	0.51
Varying by regions:	5 regions	6 regions	
Average weekly market share held by the company relative to its competitors in a given region (2018)	24.15	23.33	0.70

[23] Weekly sales and average household income amounts are converted from the local currency to US dollars.

Table 3 Descriptive Statistics

Variable	*Mean*	*Standard deviation*	*Minimum*	*Maximum*
Sales (in US$)	97,626	41,237	0.04	319,316
Post	0.349	0.477	0.000	1.000
Best Practices (BP)	0.460	0.498	0.000	1.000
Average Number of Reactions	35.542	17.200	22.500	64.857
Average Times Seen	513.223	147.846	318.000	743.000
# of Nearby Stores	16.412	22.053	0.000	78.000
Proximity to BP Stores	161.818	28.107	23.425	195.416
Prior Performance in (US$/employee)	5,659	2,888	1,962	65,747
Trust in Regional Manager Competence	-0.081	0.542	-1.193	0.925

Notes: *Average Number of Reactions* and *Average Times Seen* measure the number of reactions ("likes", "hearts", etc.) to a BP post and the number of users who saw a BP post within 2 weeks of the initial time when the post was entered on the ESN, averaged across all BP posts a focal store was exposed to over the treatment period. *Proximity to BP stores* equals 200 minus the average physical distance (in kilometres) between a focal store and the BP stores in the region. *# of Nearby Stores* measures the number of same-company stores within a 10-kilometer radius of the focal store. *Prior Performance* is the average sales ($) per employee for a focal store over the pre-intervention period. *Trust in Regional Manager Competence* is principal component factor generated by the three pre-intervention survey questions asking about employees' trust in regional manager's competence.
The number of observations is 32,564 (store-weeks), except (a) *Average Number of Reactions* and *Average Times Seen* are only generated for the 14,993 treatment store-weeks; (b) the *Proximity to BP Stores* measure is only generated for the 14,993 treatment store-weeks; (c) the number of observations for *# of Nearby Stores* is slightly lower (32,509), as we lacked location information for several stores and (d) the number of observations for *prior performance* is slightly lower (32,478), as we lacked accurate information on the number of employees for several stores.

Table 4 Does a Best Practices Initiative Redirect User Activities on the ESN?

	Dependent variable: Posts made in the Regional Groups	Dependent variable: % Posts made in the Regional Groups	Dependent variable: Comments made in the Regional Groups	Dependent variable: % Comments made in the Regional Groups
Post	-0.0095	-0.0020	-0.0128**	-0.0032***
	(-1.17)	(-1.06)	(-2.13)	(-2.65)
BP × Post	0.0710***	0.0108***	0.0402***	0.0094***
	(4.25)	(4.58)	(3.16)	(4.02)
N	27,318	27,318	27,318	27,318
Adj R^2	0.029	0.027	0.018	0.015

Note: All regressions include store, week, and store-time-trend fixed effects. Robust t-statistics in parentheses: *, **, and *** denote significance at a 0.10, a 0.05, and a 0.01 level. Standard errors are clustered by store. *Post*=1 if the week is or comes after the first week of the best practices initiative (August 26, 2019). *BP*=1 if the store is a treatment store. *Time* is the number of weeks relative to the first week of the initiative (-52 to +22). In Column 1 (3), the dependent variable measures the total number of posts (comments) made by employees in a store-week in their respective regional groups on the ESN. In Column 2 (4), the dependent variable measures the percentage of posts (comments) made by employees in a store-week in their respective regional groups relative to the total number of posts (comments) made by these employees on the ESN.

Table 5 Hypothesis 1: Does a Best Practices Initiative on an ESN Improve Financial Performance?

	Dependent variable:
	Ln(Sales)
Post	-0.0125
	(-0.70)
BP × Post	0.0034
	(0.27)
Post × Time	-0.0006
	(-0.85)
BP × Post × Time	0.0020**
	(2.48)
Adj R^2	0.094

Note: Sample size is 32,564 (store-weeks). The regression includes store, week, and store-time-trend fixed effects. Robust t-statistics in parentheses: *, **, and *** denote significance at a 0.10, a 0.05, and a 0.01 level. Standard errors are clustered by store. *Post*=1 if the week is or comes after the first week of the best practices initiative (August 26, 2019). *BP*=1 if the store is a treatment store. *Time* is the number of weeks relative to the first week of the initiative (-52 to +22).

Table 6 OLS Regressions Showing the Moderating Effects of Post Popularity on the Performance Effects of an ESN Best Practices Initiative

	Dependent variable: Ln(Sales)				
	Less Popular Posts	More Popular Posts	Moderator: Popular Posts_High	Moderator: Average Number of Reactions	Moderator: Average Times Seen
Post	-0.0059	-0.0052	-0.0126	-0.0125	-0.0126
	(-0.28)	(-0.24)	(-0.70)	(-0.70)	(-0.70)
BP × Post	0.0087	-0.0026	0.0087	0.0110	0.0085
	(0.80)	(-0.14)	(0.80)	(0.48)	(0.23)
Post × Time	0.0001	-0.0011	-0.0006	-0.0006	-0.0006
	(0.06)	(-1.28)	(-0.84)	(-0.84)	(-0.84)
BP × Post × Time	-0.00003	0.0044***	-0.00003	-0.0027*	-0.0057***
	(-0.04)	(3.98)	(-0.03)	(-1.86)	(-2.88)
BP × Moderator × Post			-0.0113	-0.0021	-0.0001
			(-0.62)	(-0.28)	(-0.12)
BP × Post × Moderator × Time			0.0044***	0.0013***	0.0002***
			(3.81)	(3.48)	(3.80)
N	25,663	24,472	32,564	32,564	32,564
Adj. R^2	0.106	0.089	0.094	0.094	0.094

Note: All regressions include store, week, and store-time-trend fixed effects. Robust t-statistics in parentheses: *, **, and *** denote significance at a 0.10, a 0.05, and a 0.01 level. Standard errors are clustered by store. *Post*=1 if the week is or comes after the first week of the best practices initiative (August 26, 2019). *BP*=1 if the store is a treatment store. *Time* is the number of weeks relative to the first week of the initiative (-52 to +22). Column 1 (2) is the subsample for which post popularity (as measured by the average number of reactions or times seen for the BP posts) is lower than or equal to (higher than) the sample median. Measuring post popularity by either the average number of reactions or average times seen for the BP posts results in the same partition of the full sample as these two measures are highly correlated. In Column 3, the moderator *Popular_High*=1 when post popularity is higher than the sample median. In Columns 4 and 5, the moderator is Average Number of Reactions or the Average Times Seen scaled by 100 (see Table 3's footnote for variable definition). Post popularity measures are only generated for the treatment stores. In each subsample analysis (Columns 1 and 2), we included all control stores. In the full sample analyses (Columns 3, 4, and 5), the moderator is set to zero for all control stores. As control stores have no variation in the moderator measures, *Post x Moderator*, and *Post x Moderator x Time* are dropped from these regressions due to multi-collinearity between all the variables.

Table 7 Hypothesis 2: OLS Regressions Showing the Moderating Effects of Prior Exposure to Knowledge on the Performance Effects of an ESN Best Practices Initiative

	Dependent variable: Ln(Sales)			
	Fewer Nearby Stores	More Nearby Stores	Moderator: # Nearby Stores_High	Moderator: # Nearby Stores
Post	-0.0219***	0.0002	-0.0057	-0.0170
	(-2.65)	(0.00)	(-0.37)	(-0.82)
BP × Post	-0.0103	0.0217	-0.0102	0.0082
	(-0.74)	(0.97)	(-0.73)	(0.60)
Post × Time	-0.0006	-0.0007	-0.0023***	0.0001
	(-0.67)	(-0.51)	(-2.72)	(0.14)
BP × Post × Time	0.0029***	0.0008	0.0029***	0.0007
	(2.61)	(0.71)	(2.61)	(0.83)
Post × Moderator			-0.0159	-0.0202
			(-0.83)	(-1.04)
BP × Post × Moderator			0.0319	0.0193
			(1.21)	(0.86)
Post × Moderator × Time			0.0039***	0.0034***
			(3.48)	(3.01)
BP × Post × Moderator × Time			-0.0021	-0.0016
			(-1.29)	(-1.26)
N	18,655	13,854	32,509	32,509
Adj. R^2	0.114	0.078	0.094	0.095

Note: All regressions include store, week, and store-time-trend fixed effects. Robust t-statistics in parentheses: *, **, and *** denote significance at a 0.10, a 0.05, and a 0.01 level. Standard errors are clustered by store. *Post*=1 if the week is or comes after the first week of the best practices initiative (August 26, 2019). *BP*=1 if the store is a treatment store. *Time* is the number of weeks relative to the first week of the initiative (-52 to +22). *# of Nearby Stores* is the number of same-company stores within a 10km-radius of the focus store. Column 1 (2) is the subsample for which *# of Nearby Stores* is lower than or equal to (higher than) the sample median. In Column 3, *# Nearby Stores_High*=1 when *# of Nearby Stores* is higher than the sample median. In Column 4, the moderator "# Nearby Stores" is a z-score transformation of the raw measure *# of Nearby Stores* (i.e. the raw measure is mean-centered and then divided by the standard error). The number of observations is slightly lower than the full sample size in the other tables because a few stores have missing information on their location.

Table 8 Hypothesis 3: OLS Regressions Showing the Moderating Effects of Geographic Proximity between the Focal Store and the BP Stores on the Performance Effects of an ESN Best Practices Initiative

	Dependent variable: Ln(Sales)			
	Less Proximity to BP Stores	Greater Proximity to BP Stores	Moderator: Proximity to BP Stores_High	Moderator: Proximity to BP Stores
Post	-0.0022	-0.0090	-0.0126	-0.0126
	(-0.10)	(-0.41)	(-0.70)	(-0.70)
BP × Post	-0.0077	0.0149	-0.0076	0.0035
	(-0.54)	(0.93)	(-0.54)	(0.28)
Post × Time	-0.0002	-0.0008	-0.0006	-0.0006
	(-0.28)	(-0.94)	(-0.85)	(-0.85)
BP × Post × Time	0.0007	0.0034***	0.0007	0.0020**
	(0.62)	(3.81)	(0.62)	(2.50)
BP × Moderator × Post			0.0225	-0.0089
			(1.32)	(-0.92)
BP × Post × Moderator × Time			0.0027**	0.0024***
			(2.36)	(4.30)
N	25,196	24,939	32,564	32,564
Adj. R^2	0.100	0.093	0.094	0.094

Note: All regressions include store, week, and store-time-trend fixed effects. Robust t-statistics in parentheses: *, **, and *** denote significance at a 0.10, a 0.05, and a 0.01 level. Standard errors are clustered by store. *Post*=1 if the week is or comes after the first week of the best practices initiative (August 26, 2019). *BP*=1 if the store is a treatment store. *Time* is the number of weeks relative to the first week of the initiative (-52 to +22). *Proximity to BP stores* equals 200 minus the average physical distance (in kilometres) between a focal store and the BP stores in the region. Column 1 (2) is the subsample for which *Proximity to BP stores* is lower than or equal to (higher than) the sample median. In Column 3, *Proximity to BP stores_High*=1 when the store's *Proximity to BP stores* is higher than the sample median. In Column 4, the moderator "Proximity to BP stores" is a z-score transformation of the raw measure *Proximity to BP store*s (i.e. the raw measure is mean-centered and then divided by the standard error). *Proximity to BP stores* is only generated for the treatment stores. In each subsample analysis (Columns 1 and 2), we included all control stores. In the full sample analysis (Columns 3 and 4), we set the control stores' moderator value to be equal to that of the minimum moderator value for the treatment stores. As control stores have no variation in the moderator measures, *Post x Moderator*, and *Post x Moderator x Time* are dropped from these regressions due to multi-collinearity between all the variables.

Table 9 Hypothesis 4: OLS Regressions Showing the Moderating Effects of Prior Performance on the Performance Effects of an ESN Best Practices Initiative

	Dependent variable: Ln(Sales)			
	Lower Prior Performance	Greater Prior Performance	Moderator: Prior Performance_High	Moderator: Prior Performance
Post	-0.0258**	0.0006	-0.0244	-0.0110
	(-2.23)	(0.02)	(-1.62)	(-0.61)
BP × Post	0.0230	-0.0163	0.0230	0.0001
	(1.18)	(-1.05)	(1.18)	(0.01)
Post × Time	-0.0018	0.0005	-0.0017	-0.0007
	(-1.47)	(0.58)	(-1.63)	(-0.90)
BP × Post × Time	0.0034***	0.0005	0.0034***	0.0018**
	(2.93)	(0.48)	(2.94)	(2.15)
Post × Moderator			0.0236	0.0076
			(1.31)	(1.02)
BP × Post × Moderator			-0.0392	-0.0444
			(-1.58)	(-1.65)
Post × Moderator × Time			0.0021*	0.0010
			(1.87)	(1.09)
BP × Post × Moderator × Time			-0.0028*	-0.0063*
			(-1.76)	(-1.95)
N	16,461	16,017	32,478	32,478
Adj. R^2	0.087	0.104	0.094	0.097

Note: All regressions include store, week, and store-time-trend fixed effects. Robust t-statistics in parentheses: *, **, and *** denote significance at a 0.10, a 0.05, and a 0.01 level. Standard errors are clustered by store. *Post*=1 if the week is or comes after the first week of the best practices initiative (August 26, 2019). *BP*=1 if the store is a treatment store. *Time* is the number of weeks relative to the first week of the initiative (-52 to +22). *Prior Performance* is sales scaled by the number of employees in a store-week averaged over the pre-intervention period. Column 1 (2) is the subsample for which *Prior Performance* is lower than or equal to (higher than) the sample median. In Column 3, *Prior Performance_High*=1 when *Prior Performance* is higher than the sample median. In Column 4, the moderator "Prior Performance" is a z-score transformation of the raw measure *Prior Performance* (i.e. the raw measure is mean-centered and then divided by the standard error). The number of observations is slightly lower than the full sample size in the other tables because data on number of employees are not available for four stores.

Table 10 Hypothesis 5: OLS Regressions Showing the Moderating Effects of Trust in Regional Manager Competence on the Performance Effects of an ESN Best Practices Initiative

	Dependent variable: Ln(Sales)			
	Lower Trust in Regional Manager Competence	Greater Trust in Regional Manager Competence	Moderator: Trust in Regional Manager Competence _High	Moderator: Trust in Regional Manager Competence
Post	-0.0231	-0.0052	-0.0101	-0.0161
	(-1.62)	(-0.18)	(-0.60)	(-0.88)
BP × Post	-0.0115	0.0132	-0.0115	0.0061
	(-0.49)	(0.92)	(-0.49)	(0.46)
Post × Time	-0.0006	-0.0006	-0.0010	-0.0005
	(-0.50)	(-0.65)	(-1.00)	(-0.62)
BP × Post × Time	0.0043***	0.0004	0.0043***	0.0018**
	(3.43)	(0.43)	(3.43)	(2.23)
Post × Moderator			-0.0044	-0.0199**
			(-0.25)	(-2.50)
BP × Post × Moderator			0.0246	0.0237**
			(0.90)	(2.51)
Post × Moderator × Time			0.0006	0.0009
			(0.55)	(1.50)
BP × Post × Moderator × Time			-0.0039**	-0.0008
			(-2.38)	(-1.12)
N	13,730	18,834	32,564	32,564
Adj. R^2	0.079	0.11	0.094	0.094

Note: All regressions include store, week, and store-time-trend fixed effects. Robust t-statistics in parentheses: *, **, and *** denote significance at a 0.10, a 0.05, and a 0.01 level. Standard errors are clustered by store. *Post*=1 if the week is or comes after the first week of the best practices initiative (August 26, 2019). *BP*=1 if the store is a treatment store. *Time* is the number of weeks relative to the first week of the initiative (-52 to +22). *Trust in Regional Manager Competence* is a measure of the level of trust employees had in the competence of the regional manager as expressed in their responses to the pre-intervention survey. Column 1 (2) is the subsample for which *Trust in Regional Manager Competence* is lower than (higher than or equal to) the sample median. In Column 3, *Trust in Competence_High*=1 when *Trust in Regional Manager Competence* is higher than or equal to the sample median. In Column 4, the moderator "Trust in Regional Manager Competence" is a z-score transformation of the raw measure *Trust in Regional Manager Competence* (i.e. the raw measure is mean-centered and then divided by the standard error).

Appendix 1: *Ex Ante* Power Analyses

Treatment is applied to the treated group starting on August 15, 2019. Below, we consider effects of the treatment on overall sales.

The data are simulated using fixed effects of stores and week-of-the-year as well as taking into account trends for each individual store. For store i and week t, the log-sales are estimated as follows:

$$y_{it} = \alpha_i + \beta_i t + \gamma_t + \delta D_{it} + \epsilon_{it},$$

where $D_{it} = 1$ if store i is treated at time t (i.e., t is any week post August 15, 2019); δ is a treatment effect (on a log scale); α_i and β_i are, respectively, fixed effects and trend of store i; and γ_t is a fixed effect of week t.

Note that, before the simulation, we exclude "volatile" weeks from the data. The following table specifies the holiday week numbers that are excluded in the simulation from 2015 to 2019.

Holidays and vacation periods	**Year**				
	2015	**2016**	**2017**	**2018**	**2019**
New Year's Day	1	53	52	1	1
Other local holidays	14, 15, 20, 22	12, 13, 18, 20	15, 16, 20, 21, 23	13, 15, 19, 20, 21	16, 17, 20, 22, 24
Summer break	25-33	25-33	25-33	25-33	26-33
Christmas break	52-53	51-52	51-52	51-52	51-53

After simulating hypothetical data for 2018 and 2019, we perform power analysis by varying the size of the treatment effect. Below are the results obtained for each outcome considered.

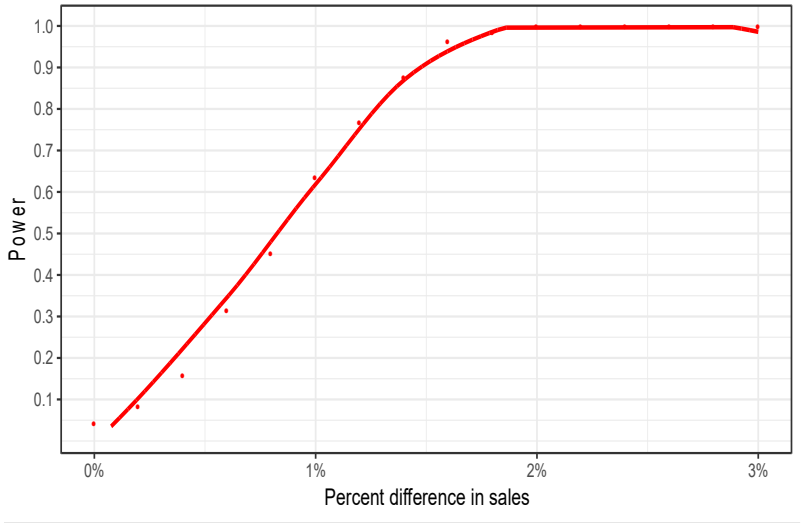

(Above) Estimated Power for All Sales

Appendix 2: Examples of Best Practices Posts[24]

POST 1

BEST PRACTICES ON BEVERAGES

Picture:

The picture of this post showed a clean and organized aisle, tightly filled (from the floor to the roof and from side to side) with different types of beverages

Text included with the post:

- Keeping Shelves Filled

Independent of turnover at the store, this department should always be filled up. Nothing is as ugly as a poor soda-department, and likewise nothing is as fantastic as a well filled-up soda-department either. There is no special risk for waste with beverages. At [STORE NAME], there is only a very small storage for products, but beverages are the big exception. Here, we should always be able to fill up such that the department is bursting with products. It is recommended to have one person that has a little extra responsibility for restocking products here.

Even if [STORE NAME] experiences a trend where more customers want to explore the range of beer-assortment, it is important to have an assortment of beer that is "correct" for the customer base. [We] use numbers and customer insights to adjust the department after the local population's preferences.

POST 2

BEST PRACTICES ON BREAD

Picture:

A picture was inserted showing an open area for bread including tall shelves arranged in an L-shaped with a lower-level display in the center.

Text included with the post:

[STORE NAME] rebuilt the bread department a couple of weeks ago. Earlier they had 2 bread fronts that made the customers have to walk around the entire disk to see the entire assortment. Now the department is more open and the customers can see more of the assortment when they walk toward the bread department.

[24] With the exception of the picture in Figure 1, the company's management requested that we do not share the pictures displayed in the best-practices posts for confidentiality purposes.

POST 3

BEST PRACTICES ON DRY GOODS

Picture:

This picture showed a highly organized aisle next to the register tightly filled with a wide variety of products such as lozenges, chewing gum, snacks, etc. The picture resembled a lengthy tightly packed duty free counter at an airport.

Text included with the post:

See Potential at the Checkout Zone

Here there are goods with high gross margins, that could get lost if you do not prioritize this space.

- [STORE NAME] fights for the top position, with revenue of over [$$$$] only on gum/lozenges at the register zone so far in 2019. This accounts for lots of gross [sales].
- Achieve more sales area by the register by sharing/dividing the table. Smaller stores have lots to gain by taking advantage of this area.
- This is a picture of register 2, impulse fruits & vegetables are displayed in register 1

Avoid Static End Aisles

- Try to have a max. time horizon of 2-3 [weeks] at the end aisle
- Own brands have a static in/out price, but keep making changes with other products to surprise the customers, increase gross [sales], and reduce waste.
- Spend time acquiring knowledge about which items can be sold extensively.

POST 4

BEST PRACTICES ON FRESH PRODUCTS

Picture:

The picture of this post showed two individuals (presumably the franchise owner/store manager and the department manager of fresh goods) holding in both of their hands two ready-made dinners each, behind a cold display counter where different packets of ready-made dinners were displayed cross-sectionally in a highly organized fashion.

Text included with the post:

The store has so far this year a +5.4% growth on fresh goods. The franchise owner with the person in charge of fresh goods plan the weekly disks together, weekday and weekend. They have had a special focus on SRDs (simple, ready dinners under [$$$ price]) and ensured that there are good and simple exposures of high rolling SRD goods in the counters. The customers love it!

POST 5

BEST PRACTICES ON FRUITS AND VEGETABLES

Picture:

Seven pictures were included in the album featuring highly organized fruit-and-vegetable displays following the guidelines described in the text of the post (see below). Each of the pictures could be accessed through a click, and each had an explanation of how that display followed the principles shared in the main post.

Text included with the post:

Together with franchise owner XX and F&V responsible XX… we took some simple steps to increase growth in the department. We rebuilt the department in week 32 and positioned it according to the following principles:

1. The right item in the right place!
2. Sell more of what you sell a lot of! In other words, bring up the volume and growth of bestsellers!
3. Counter: one price per whole counter! Max two products, two prices per entire counter!
4. Priority goods on counter: High rolling goods on all counters!

   Priority products we recommend for the counters:

   - Avocado 2pk
   - Avocado ripe single
   - Mango 2pk
   - Mango single
   - Cherry tomatoes: our best mini-plum tomatoes [XX Name of tomatoes XX]
   - Apple pink lady or current royal gala
   - Sugar peas in finished bags (not by weight)
   - Snack carrot
   - Berries: blueberries, raspberries, strawberries
   - Pointed peppers
   - Sweet potatoes
   - Season: berries, plums, cherries and more.