This article was downloaded by: [216.73.217.54] On: 26 June 2026, At: 07:57
Publisher: Institute for Operations Research and the Management Sciences (INFORMS)
INFORMS is located in Maryland, USA

Information Systems Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Research Spotlights

To cite this article:
(2026) Research Spotlights. Information Systems Research 37(2):iv-xii. https://doi.org/10.1287/isre.2026.resspot.v37.n2

Full terms and conditions of use: https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2026, INFORMS

Please scroll down for article—it is on subsequent pages

With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.
For more information on INFORMS, its publications, membership, or meetings visit http://www.informs.org

https://pubsonline.informs.org/journal/

INFORMATION SYSTEMS RESEARCH
Vol. 37, No. 2, June 2026, pp. iv–xii
ISSN 1047-7047 (print), ISSN 1526-5536 (online)

Research Spotlights

https://doi.org/10.1287/isre.2026.resspot.v37.n2

Copyright: © 2026 INFORMS

Downloaded from informs.org by [216.73.217.54] on 26 June 2026, at 07:57 . For personal use only, all rights reserved.

Toward Artificial Intelligence Compliance: Impacts and Mechanisms of Performance Feedback (p. 667)
Shaobo Wei, Yuanyuan Zhang, John Qi Dong

As organizations increasingly adopt artificial intelligence (AI) to enhance performance, ensuring that employees use AI in compliance with organizational policies becomes crucial for realizing its full value. However, employees' AI compliance is not guaranteed and can vary based on how their AI use is managed. This study offers timely and actionable insights into how performance feedback—both positive and negative—influences employees' AI compliance, and how these effects vary with AI identity. Drawing on feedback intervention theory, we conduct a longitudinal field study and a randomized experiment and find that positive performance feedback promotes AI compliance, whereas negative performance feedback reduces it. Importantly, employees with high AI identity respond more strongly to both types of performance feedback. Our findings further uncover distinct underlying mechanisms—task-motivation, task-learning, and meta-cognitive processes—that channel the effects of positive and negative performance feedback on AI compliance. Taken together, organizations should tailor performance feedback as part of AI governance by considering employees' AI identity. Positive reinforcement of AI compliance is especially effective for employees with high AI identity, whereas cautions are needed when delivering negative performance feedback to avoid undermining AI compliance. Policy guidelines should support identity-sensitive performance feedback in practice.

Sponsored Tasks and Solver Participation in Crowdsourcing Contests (p. 693)
Jiahui Mo, Sumit Sarkar, Jianqing Chen

Crowdsourcing platforms enable firms to solicit solutions from external contributors (solvers). As the popularity of these platforms grows, many similar tasks compete for solver attention, leading to insufficient participation for some contests. Platforms offer sponsorship options—paid features that highlight tasks to increase visibility. This study investigates whether such sponsorship effectively attracts more solvers. Sponsorship remains underexplored in crowdsourcing contests, despite its widespread use in many crowdsourcing platforms. Given the unique competitive dynamics of contests—distinct from traditional e-commerce—the effectiveness of sponsorship is difficult to anticipate. We conduct both task- and task-day level analyses. We find that sponsorship increases participation only when paired with high prize amounts; even then, the increase is driven mainly by low-ability solvers. The rise in total submissions stems from more participants rather than increased quantity per solver. We identify anticipated competition as a novel mechanism through which sponsorship can deter capable solvers, complicating the net effect. This study contributes to both the crowdsourcing and sponsorship literatures. Unlike prior work that primarily emphasizes the positive effects of sponsorship, we reveal its nuanced consequences in all-pay auction environments. Platforms and firms should be cautious when using sponsorship, as solvers respond differently depending on their ability.

Game for Brainstorm: The Impact of a Badge System on Knowledge Sharing (p. 716)
Lei Wang, Yifan Zhang, Yi-Jen (Ian) Ho

Gamification systems such as badge rewards are widely used to encourage user engagement, yet their effectiveness depends heavily on design. This study investigates how badge volume, variety, and valence—the 3Vs—influence knowledge sharing on platforms such as Stack Overflow. Using a structural hidden Markov model with a copula correction for endogeneity, we uncover that badge volume and valence significantly increase both the quantity and quality of user contributions, particularly among inactive and experienced users. However, excessive variety in badge types reduces long-term engagement, suggesting that diversification may dilute motivational focus. Counterfactual simulations reveal that reducing the difficulty of earning answer badges ("volumizing") enhances contributions, whereas altering the rarity of high-valence badges (e.g., gold) often backfires. These insights highlight the importance of aligning gamification mechanics with user psychology and engagement trajectories. Practically, platforms can personalize badge offerings, offering easy wins for newcomers and prestige rewards for experienced users. Aligning badge incentives with

Downloaded from informs.org by [216.73.217.54] on 26 June 2026, at 07:57 . For personal use only, all rights reserved.

content goals (e.g., Python questions) and visually showcasing rare badge collections can further deepen engagement. Our findings offer actionable guidance for platform designers to refine gamification systems that balance challenge, recognition, and motivation.

Geographical Pattern of Online Word of Mouth: How Offline Environment Influences Online Sharing (p. 736)
Tianshu Sun, Yanhao "Max" Wei, Joseph Golden

A key feature of e-commerce is the access to a wide range of geographical areas; e-commerce platforms spend much effort customizing their strategies across geographical areas. Meanwhile, online word of mouth (WOM) is becoming an increasingly important driver for product sales on e-commerce platforms. This study investigates the geographical pattern of online WOM. Specifically, we examine whether and how customers' local environment influences the generation and direction of online WOM. Leveraging a unique research design, we measure the online WOM at a U.S. national e-commerce platform by tracking the WOM referral decisions of customers. Whereas digital technologies have enabled customers to share without physical boundaries, we find that location still plays significant roles in the generation and direction of online WOM. First, a customer's offline social environment (e.g., friend visiting frequency, neighbor interactions) significantly explains the generation of online referrals. Second, online referrals are largely bounded locally, and referrals in areas with more local social interactions are more likely to stay local. Third, even when referrals travel far, they are more likely to point to destinations socially similar to the origins. We derive important implications for firms in allocating resources across geographical areas and in using referrals to identify high-potential geographical markets.

Third-Party Software Development Kit Utilization and Mobile App Market Performance (p. 761)
Yu Xia, Hailiang Chen, Yulin Fang

In the highly competitive mobile market, third-party vendors located outside the purview of hosting mobile platforms are becoming major suppliers of functional tool kits for mobile app development and innovation. However, app developers face challenges in navigating the use of third-party software development kits (SDKs) to enhance app performance. Drawing on the platform ecosystem literature, this study examines the roles of third-party SDKs, platform owners, and mobile app developers from a boundary-spanning perspective, emphasizing their respective contributions to shaping the functional value of third-party SDKs as boundary objects. By analyzing a longitudinal data set of 335,958 apps from the Apple App Store and Google Play, we find that integrating more third-party SDKs can significantly boost daily active users, especially when tool-type SDKs are utilized. However, this positive effect diminishes following platform updates that redefine platform boundaries and as developers gain more platform-specific experience. This study highlights the importance of balanced governance in platform ecosystems to sustain market growth, offering actionable insights for developers, policymakers, and industry leaders in shaping the future of mobile app markets.

Quality Control for Crowd Workers and for Language Models: A Framework for Free-Text Response Evaluation with No Ground Truth (p. 782)
Inbal Yahav, Anat Goldstein, Tomer Geva, Shahar Meir, Onn Shehory

As businesses increasingly rely on large language models (LLMs) for tasks such as customer service and information retrieval, ensuring the accuracy of their responses is a critical challenge. Traditional verification is costly, slow, and often requires scarce domain experts. We introduce the automated quality evaluation based on textual responses (AQER) framework, a novel, cost-effective method to assess the correctness of free-text answers from both LLMs and human workers without needing preexisting correct answers. AQER works by intelligently aggregating multiple responses to the same question, leveraging the wisdom of the crowd to create a reliable synthetic correct answer, followed by an iterative procedure that accounts for response quality cues. AQER obtains state-of-the-art performance compared with existing automated response evaluation baselines. For managers AQER offers a scalable, data-driven method to (i) evaluate and select the best performing LLMs for specific organizational needs and use cases, (ii) continuously monitor artificial intelligence (AI) performance to ensure reliability and accountability across different model versions, and (iii) manage the quality of crowd workers essential for high-quality AI training and validation. AQER, thus, offers a robust mechanism for improving model performance and mitigating the significant financial and reputational risks associated with deploying untrustworthy generative AI technologies.

Impacts of Reducing Visibility of Friends' Liked Content on User Content Engagement Across Newsfeed Channels (p. 805)
Xiaohui Zhang, Qinglai He, Zhongju Zhang

Nowadays, online platforms constantly adjust their newsfeeds to boost engagement, often emphasizing nonsocial channels such as algorithmic recommendations over social ones such as user networks. However, ignoring the interactions between these channels can have unintended consequences. This study delves into this by assessing the impacts of a policy change that reduced the visibility of friends' liked content on a major discussion platform. We found this led to an overall decrease in users' content engagement. Although users

Downloaded from informs.org by [216.73.217.54] on 26 June 2026, at 07:57 . For personal use only, all rights reserved.

engaged more with their friends' original posts and trending topics, their interaction with nonsocial content such as algorithmic recommendations decreased. This reveals a key insight: social channels act as substitutes for one another but are complementary to nonsocial, algorithmic channels. Vibrant social activity is crucial for driving traffic to other parts of a platform. Crucially, the change also made users' content engagement less diverse. Friends' liked content is a key source of exposure to niche topics. For platform operators, this means that sidelining social features in favor of algorithmic feeds can backfire, reducing overall engagement and diversity. For policymakers concerned about online echo chambers, our findings suggest that content shared through extended social networks is vital for promoting a wider range of content.

Artificial Intelligence-Powered Digital Streamers in Online Retail: Empirical Insights and Design Strategies from Experiments (p. 824)
Yahui Liu, Lei Wang, Shuai Yang, Yanwen Wang

As artificial intelligence (AI)-powered digital streamers gain popularity in live commerce, online retailers face critical questions about the actual business value of their operations. This study offers timely, evidence-based insights into the economic impact and optimal design of digital streamers. Although current designs do not significantly improve sales over no live streaming, incorporating behavioral realism—especially enhanced real-time question and answer (Q&A)—can boost sales by 25%, making digital streamers as effective as human hosts. Visual upgrades and human-like voices also help but to a lesser degree. Importantly, not all AI-driven enhancements deliver immediate returns, and imitating human scripts does not guarantee success. Retailers should focus on dynamic human-AI interaction features that drive engagement and trust, such as real-time Q&A and interactive giveaways. Designers are encouraged to integrate multiple realism features to maximize effectiveness while managing cost and scalability. These findings offer actionable guidance for retailers and platform designers seeking to leverage AI effectively and cost efficiently in live streaming commerce.

Is Prevention Better Than Cure? Effects of Cyber Risk Disclosures on Shareholder Response to Breaches (p. 842)
Rui Cao, Moksh Matta, Hasan Cavusoglu, Arslan Aziz, Özüm Kafaee

As digitalization increases firms' exposure to cyber risks, corporate disclosures about how these risks are managed are becoming more common and influential. This study examines 1,912 breach incidents affecting public companies to understand how shareholder reactions differ depending on the type of cyber risk strategies disclosed. We find that, although breaches generally lead to stock price declines, firms that previously disclosed *preventive* strategies, such as efforts to avoid breaches, experience significantly smaller losses in market value. Conversely, disclosing *mitigative* strategies, focused on damage control after a breach, amplifies the negative impact. These effects arise from shareholders' loss aversion: They respond more favorably to firms perceived as trying to prevent harm rather than simply reacting to it. These findings suggest that managers should focus cyber risk disclosures on credible, prevention-oriented strategies to build investor confidence and minimize financial fallout after a breach. Additionally, our findings advise against using cyber risk disclosures as tools for impression management. Managers should ensure these disclosures accurately reflect the firm's cyber risk management practices, as failing to do so can undermine the economic benefits of emphasizing preventive strategies.

CEO Human Capital and Digital Product Innovation: A Dynamic Managerial Capabilities Perspective (p. 863)
Colin Schulz, David Bendig, Johannes Kriebel, Kathrin Haubner, Stav Fainshmidt

How does the human capital of chief executive officers (CEOs) drive digital product innovation in manufacturing firms? Analyzing data from 216 U.S. firms and over 8,000 new product announcements, this study shows that technological and business knowledge can both enable and inhibit digital product innovation, depending on the external environment. In stable settings, tech-savvy CEOs drive digital product innovation, whereas business-savvy CEOs focus elsewhere. In dynamic environments, these effects reverse. These findings offer a contextual view on the ongoing debate about the value of technological versus business expertise in top management, suggesting that neither type of knowledge is universally beneficial. We also find that CEOs serve as distinct innovation catalysts beyond their top management teams. Follow-up interviews reveal the diverse strategies they use to initiate, develop, and implement digital product innovation in established firms. Based on these findings, the study offers guidance for boards of directors on aligning CEO selection with environmental demands and for CEOs seeking to expand their knowledge base to more effectively foster digital product innovation under varying conditions.

Transferring, Eliminating, or Both? The Strategic Impact of Resale Platforms and Consumer Learning on Coping with Consumer Uncertainty (p. 886)
Yongjian Li, Yusheng Wang, Xin (Robert) Luo, Xuanming Bai

For sellers of experience goods, consumer uncertainty threatens profits. Resale platforms and enhanced

Downloaded from informs.org by [216.73.217.54] on 26 June 2026, at 07:57 . For personal use only, all rights reserved.

consumer learning are popular tools to manage this uncertainty, but their impacts differ significantly and require tailored strategies. Our research reveals critical nuances overlooked by prior studies: the resale platform's negative effect is often underestimated because of its indirect influence on consumer waiting behavior, whereas consumer learning can harm sellers of high-cost goods when consumers are impatient. Critically, the profitability and environmental outcomes depend on production costs and consumer patience. Encouraging consumer learning creates a profit-environment win–win with patient consumers, whereas resale platforms excel with impatient consumers and high costs. Combining both mechanisms benefits sellers only when consumers are patient and costs are moderate ("1 + 1 > 1"). Key strategic imperatives emerge: (i) resale behavior: encourage selectively (beneficial mainly for high-cost goods), minimizing resale discounts; (ii) learning behavior: nurture consumer patience and avoid rushing purchases, especially for premium goods; and (iii) product upgrades: avoid frequent, minor updates. Sellers must move beyond one-size-fits-all approaches, aligning their uncertainty management (resale, learning, upgrades) with product economics and consumer patience to optimize profits and sustainability.

A General-Purpose IT Intervention to Improve Human Decision Making, Strengthen Passwords, and Reduce Receptivity to Misinformation (p. 904)
Kevin Andrew Harmon, Eric A. Walden

Latency—the delay between a user's action and a system's response—is often viewed as a problem to be eliminated. Yet our research shows that latency is not always harmful and can even improve decision making under the right conditions. Across two sets of studies, we find that latency enhances decision accuracy when it occurs after task-relevant information has been presented but before a decision is made. In these cases, the pause functions as a "mental speed bump," prompting users to shift from quick, intuitive judgments to slower, more reflective thinking. However, when latency occurs without relevant content, performance declines, likely because of mind wandering and disengagement. For practitioners, this means that latency can be deliberately integrated into system design to encourage deeper cognitive processing in high-stakes or complex tasks, such as financial decisions, medical diagnostics, or security-critical actions. For policymakers, the findings suggest a shift in focus: rather than universally minimizing latency, standards and best practices could account for situations in which strategic delays benefit user outcomes. By reframing latency as a design variable rather than a nuisance, organizations can leverage it to support more accurate, thoughtful decision making, improving performance in contexts in which getting the right answer matters most.

Review Manipulation and Filtering on Digital Platforms (p. 927)
Xian Gu, Jingcun Cao, Yulin Fang

Manipulated consumer reviews are a growing concern for digital platforms, undermining trust and distorting market outcomes. This study analyzes data from the Apple App Store to assess how both positive and negative manipulated reviews—later filtered by the platform—affect app sales rankings. Surprisingly, both one-star and five-star manipulated reviews initially boost app rankings, even when one-star reviews are intended to harm competitors. These effects can persist for weeks and take up to six months to reverse through platform filtering. Using text analysis, we find that negative manipulated reviews are linguistically more distinguishable from organic ones than are positive manipulated reviews, making them easier for consumers to spot. Our results show that review volume often outweighs valence in influencing consumer behavior, and that manipulated reviews have stronger effects on free apps, gaming apps, and apps from large developers. These findings underscore the urgency for platform managers to invest in faster, more accurate filtering systems, and highlight the need for policymakers to strengthen governance mechanisms to protect marketplace integrity and consumer trust.

Programming Tasks Impact Responses to Moral Dilemmas for Novice Programmers (p. 948)
Tanya Singh, Jui Ramaprasad, Kartik K. Ganju

The rapid diffusion of programming skills across education and industry may impact how individuals consider moral dilemmas. Across a series of experiments, we show that performing even simple programming tasks shifts novice programmers' evaluation of the classic trolley problem toward utilitarian responses. After solving a programming problem, respondents are more willing to sacrifice one life to save many. This effect arises because programming induces a deliberative, rule-based cognitive style. However, the effect diminishes with greater programming experience and can be mitigated through interventions, such as time delays or moral nudges. These findings highlight that organizations training employees in coding should be aware that programming tasks may temporarily alter moral reasoning, potentially influencing judgments in ethically charged contexts (e.g., product design, risk management, or AI development). Incorporating reflective cooling-off periods or explicit ethical reinforcement may reduce bias toward utilitarian reasoning. As programming becomes a baseline skill across the workforce, its cognitive spillovers could shape societal attitudes toward contested moral dilemmas. These include ethical trade-offs in settings such as autonomous vehicles and artificial intelligence systems.

Downloaded from informs.org by [216.73.217.54] on 26 June 2026, at 07:57 . For personal use only, all rights reserved.

Downloaded from informs.org by [216.73.217.54] on 26 June 2026, at 07:57 . For personal use only, all rights reserved.

can leverage OHCs as a strategic tool to advance equitable, high-quality care.

Beyond Pairwise Network Interactions: Implications for Information Centrality (p. 1068)
Sandro Claudio Lera, Yan Leng

Organizations and policymakers increasingly rely on network metrics to decide whom to inform, monitor, or support. Yet most networks treat interactions as pairs, even when the underlying activity occurs in groups—project teams, chat channels, meetings, or news articles that mention multiple firms. Collapsing groups into one-to-one links can misidentify who matters. We propose a practical alternative: Model group interactions directly as a hypergraph and compute centrality from a two-step diffusion process that captures how information moves across and within groups. The approach provides a transparent way to incorporate domain knowledge (e.g., whether people enter large or small groups first) and produces testable interpretable rankings. We evaluate the method in three settings—open-source software, a high school interaction study, and financial comentions—and find that hypergraph-based, theory-informed centrality better explains outcomes such as project success, student popularity, and same-day returns than standard graph centralities. For practice and policy, this yields more effective targeting, earlier warning signals, and improved allocation of attention and resources in collaborative work, public health, and market surveillance. We release an open-source Python package (HyperCentral) to support adoption.

Why Is the Grass Always Greener on the Other Side? Tourist Bias in Online Restaurant Ratings (p. 1094)
DaPeng Xu, Xiaoquan (Michael) Zhang, Hong Hong, Qiang Ye

An important assumption underlying online ratings' helpfulness is that rating valence can serve as a good proxy for the quality of a product or service. However, consumers with different backgrounds may have different quality perceptions of the same product or service. As one of the most common types of reviews, restaurant reviews can come from two very different groups of customers: locals and tourists. Merely showing the average rating of a restaurant by pooling all ratings may not be helpful for consumers who wish to determine whether a given restaurant is a good match for them. We find that traveler consumers are much more likely than local consumers to provide higher restaurant ratings, and they tend to attach more pictures, write shorter reviews, and use fewer cognitive words. Individuals' changes in focus and changes in evaluation processes induce such a tourist bias in ratings. Meanwhile, travel destination, restaurant cuisine authenticity, and consumers' consumption pattern changes are not responsible for the bias. Our study demonstrates the necessity of differentiating travelers' and locals' ratings/reviews and identifies possible reasons behind the tourist bias, which can help retailers, consumers, and platforms better understand, utilize, and manage digital word-of-mouth in practice.

Unveiling the Strategic Impacts of Extending Membership-Based Free Shipping Programs Beyond the Online Marketplaces (p. 1113)
Geng Sun, Huseyin Cavusoglu, Srinivasan Raghunathan

Online marketplaces are increasingly extending their membership-based free shipping (MFS) programs to include external sellers, transforming logistics capabilities into a "logistics-as-a-service" (LaaS) model. This study uses a game-theoretic approach to analyze the implications of such extensions, with Amazon as a leading example. Contrary to the widespread belief that LaaS is a stand-alone profit center, our findings reveal that the primary value of extended MFS lies in enhancing the marketplace's core business. By attracting member consumers and optimizing fulfillment operations, marketplaces can increase commission revenues and reduce per-order logistics costs—even if LaaS itself operates at a loss. Importantly, this strategy can benefit the broader marketplace ecosystem, including both internal and external sellers and member consumers, while potentially disadvantaging nonmembers. For practitioners and policymakers, these results underscore the need to rethink competitive dynamics and pricing strategies in digital platforms. Rather than viewing extended MFS solely as a logistics monetization tool, it should be seen as a strategic lever for shaping marketplace structure, seller participation, and consumer behavior.

Between Human and System Agency: Coping with Negative Incidents for Continued Effective Use of Wearables (p. 1133)
Annamina Rieder

Wearable devices hold significant promise for promoting healthy behaviors, yet they often fall short of this potential when users disengage or use them ineffectively. This research examines how individuals respond to negative incidents—such as frustrating feedback, misaligned goals, or perceived surveillance—and how these responses influence continued effective use. Effective use means interacting with the wearable in ways that help achieve health-related goals, beyond merely logging steps or checking data. Based on in-depth accounts from long-term users, the study reveals that, although some cope by re-engaging with the technology, others manipulate data, disengage, or selectively avoid features, undermining the benefits wearables are meant to deliver. Crucially, sustained effective use depends not just on motivation or usability but on the alignment between human and system agency. When wearables

Downloaded from informs.org by [216.73.217.54] on 26 June 2026, at 07:57 . For personal use only, all rights reserved.

assert their own logic too strongly or misalign with user goals, maladaptive responses are more likely. These findings offer actionable guidance for designers, healthcare providers, and policymakers: Rather than focusing solely on adoption or persuasive design, efforts should support user autonomy and recovery after setbacks. Wearables that accommodate breakdowns and empower users in the face of friction are more likely to sustain engagement and improve long-term health outcomes.

Do We Think Differently When Tapping the Screen or Clicking the Mouse? Effects of Computer Interfaces on Level of Construal (p. 1158)
Xixian Peng, Xinwei Wang, Yutong Guo, Hock Hai Teo

The touchscreen has become the de facto interface for individuals to interact with digital devices. This research investigates the implications of the profound shift from the mouse-based, nontouch interface to the touch-based interface for human thinking and behavior. Ten studies reveal that using touchscreens, as opposed to a mouse, fosters a more concrete low-level construal. Specifically, using touch interfaces on devices such as iPads and tablets prompts people to interpret actions in a detailed, hands-on manner, viewing locking a door as inserting a key into a lock rather than as a broader concept such as securing a house. Touchscreens also encourage a narrow categorization of concepts and a reliance on local, granular attributes. The concrete mindset evoked by the touch interface elevates individuals' preference for practicality over complexity, such as choosing a compact camera over an advanced one, and receptiveness to persuasive attempts with feasibility appeals. Given that mental construals fundamentally influence how we perceive and engage with the world, our findings underscore the wide-ranging impacts of touch interfaces. This research provides a theoretical foundation for exploring the broader implications of the widespread adoption of touch technology, especially as digital devices become increasingly integrated into our digitally transformed daily life worldwide.

The Creation of Immersive Experiences in Transcultural Entertainment: An Action Design Process Focused on Neural Rendering (p. 1180)
Mike Seymour, Barney Tan, Yangting Li

This study examines how artificial intelligence (AI), specifically neural facial reenactment (NFR), can address the limitations of dubbing and subtitling in adapting foreign films. Using action design research and guided by presence theory, we codeveloped and evaluated an NFR process during the English adaptation of the Polish feature film *The Champion*. Unlike conventional dubbing, which often disrupts immersion through poor lip-sync or script changes, NFR preserves the original actors' performances, aligning them with new dialogue. Independent broadcast-quality assessments confirmed technical validity, and subsequent commercial distribution on a major streaming platform demonstrated scalability and audience acceptance. From this process, we derived six design principles: avoid forced script changes, respect creative intent, minimize intrusive technology, reduce training data requirements, enable flexible audience access, and codesign with existing creative structures. These principles offer a replicable template for the responsible researching and developing of AI in information systems. For practice, the findings show that NFR can improve cultural accessibility and create new creative and technical roles rather than displacing talent. For policy, the study highlights the importance of codesign, transparency, and preserving artistic integrity when integrating AI into global cultural products.

Chat More and Contribute Better: An Empirical Study of a Knowledge-Sharing Community (p. 1202)
Xiaomeng Chen, Chris Forman, Michael Kummer

Knowledge-sharing communities face a dual challenge: they must enable efficient content generation and also provide spaces in which interpersonal bonds can form. We examine how these goals interact by studying Stack Overflow's introduction of informal chat rooms alongside its main question-and-answer (Q&A) forum. Using a difference-in-differences design, we analyze two sequential shocks: the launch of chat rooms and the later addition of automated feeds that pushed Q&A content into chat rooms, thus creating a bridging mechanism between the two spaces. We find that, on their own, the introduction of chat rooms leads to no improvement in knowledge exchange for average users in the community. However, we find evidence that chat rooms may forge a closed subgroup of users, who tend to primarily help each other by answering questions within the group, and the benefits are isolated to the subgroup. The activation of an automated feeding function that pushes all questions into the chat room expands benefits beyond the closed subgroup. Our results show that platform designers should consider not only how to foster subgroup cohesion, but also how to facilitate bridging between interactions in the subgroups and the broader community.

Strategic Drivers of Core Expansion on Software Platforms: Evidence from Apple iOS (p. 1222)
André Halckenhaeusser, Jens Forderer, Armin Heinzl, Ola Henfridsson

Software platforms rely on a core-periphery structure; platform owners provide foundational code that third-party developers extend with complementary software. This structure creates strong incentives for platform owners to cooperate with third-party developers. Yet, platform firms, like Apple, intermittently expand their core by integrating peripheral functionality—for instance, flashlight or screen-time tracking in iOS. If ecosystems thrive on cooperation, why do platform owners

Downloaded from informs.org by [216.73.217.54] on 26 June 2026, at 07:57 . For personal use only, all rights reserved.

expand into niches at the apparent expense of developers? Using a comprehensive data set of all Apple iOS core expansions from 2012 to 2020, we find that expansions systematically target underperforming niches—those with low user satisfaction, limited innovation activity, and high market concentration. Stand-alone apps tend to occur in niches with low innovation, whereas integrated core features focus on niches with low user satisfaction and high concentration. For policymakers, these insights inform regulatory discussions on platform power and the design of public interventions. For platform owners and third-party developers, the findings highlight characteristics of ecosystem niches that are vulnerable to core expansion, supporting strategic decisions regarding cooperation, competition, and innovation investment. By linking niche characteristics to core expansion strategies, this study provides actionable guidance for regulating and managing evolving platform ecosystems.

Can Crowdchecking Curb Misinformation? Evidence from Community Notes (p. 1242)
Yang Gao, Maggie Mengqing Zhang, Huaxia Rui

Misinformation spreads rapidly on social media, whereas traditional countermeasures struggle to balance effectiveness, scalability, and free expression. Many platforms are now experimenting with crowdsourced fact-checking—systems that rely on users' collective judgment to identify and annotate misleading content. This paper investigates the efficacy of such systems in curbing misinformation in the context of Community Notes, a pioneering crowdsourced fact-checking system from Twitter/X. Using a regression discontinuity design, we find that publicly displaying community notes significantly increases and accelerates the voluntary retraction of misleading tweets, demonstrating the viability of crowd-based fact-checking as an alternative to professional fact-checking and forcible content removal. The effect is primarily driven by authors' reputational concerns and social pressure when corrections are visible to the public. Our findings carry meaningful implications for practice and policy. Individuals can play an active role by contributing to crowdchecking, strengthening collective information integrity. Platforms should adopt transparent, community-based systems, like Community Notes, as scalable, less controversial alternatives to forcible content removal. Policymakers can support these initiatives through regulatory guidance that promotes transparency and accountability.

What, Why, and How: An Empiricist's Guide to Double/Debiased Machine Learning (p. 1259)
Bowen Shi, Xiaojie Mao, Mochen Yang, Bo Li

We provide an introduction to double/debiased machine learning (DML), a framework that enables effect estimation when dealing with complex, high-dimensional data. In many empirical analyses, especially in fields such as information systems, researchers face difficult choices about which control variables to include and how to model their relationships with the outcome. These modeling decisions can significantly change results, leading to uncertainty about which findings are reliable. DML offers a practical solution by combining modern machine learning with rigorous statistical inference. The idea is to let flexible ML models (such as random forests or gradient boosting) capture complex relationships among control variables while still delivering reliable estimates for the key effect of interest. DML can be applied to many familiar research designs, including standard regression with controls, instrumental variables, difference in differences, and models that incorporate ML-generated features. Empirical studies and simulations show that DML is typically more robust to misspecification than traditional regression and more reliable than earlier semiparametric methods. However, DML is not automatic—it still requires sound research design and high-quality machine learning estimation. Used thoughtfully, DML provides a powerful, flexible, and statistically grounded approach for empirical research in modern data environments.

Chief Information Security Officers on Top Management Teams: Impact on Firms' Innovation (p. 1276)
Yiwen Gao, Sunil Wattal, Jason Thatcher

The growing frequency of information security breaches and the rising importance of cybersecurity have prompted many firms to include chief information security officers (CISOs) in their top management teams (TMTs). Although CISOs are often viewed narrowly through a security-focused lens, our research shows that their inclusion in TMTs can offer a strategic advantage by significantly enhancing firm innovation. We identify three mechanisms that explain this effect: (1) reducing preventable security risks that might otherwise hinder innovation efforts; (2) enabling the adoption of innovation technologies (e.g., cloud computing, big data) that carry strategic security risks; and (3) strengthening security controls that protect intellectual property and mitigate innovation-related threats. Importantly, the CISO's background matters. Those with specialized experience—either in the same industry or with prior executive roles—have a stronger impact on driving innovation. This research illuminates how CISOs' presence on TMTs affects firms' value creation from a security risk management perspective, and provides guidance for firms seeking to hire CISOs for innovation.

Learning When Reading: Evidence from an Online Mobile Reading Platform (p. 1289)
Yuchen Liu, Lizheng Wang, Yong Tan, Yongjun Li

How do consumers listen to peers during consumption? On emerging digital platforms, real-time comments

Downloaded from informs.org by [216.73.217.54] on 26 June 2026, at 07:57 . For personal use only, all rights reserved.

embedded within content—such as books or videos—allow users to observe others' spontaneous reactions while engaging with the same material. Using data from a leading Chinese online reading platform, we develop a Bayesian learning model to quantify how these in-consumption comments help users learn book quality and make sequential chapter-by-chapter purchase decisions. Results show that not all comments are viewed equally: those reflecting plot-based insight or strong narrative engagement enhance perceived quality, whereas speculative or even purely cheerful comments may diminish it. Comment *consistency*—not just volume—plays a critical role in sustaining user engagement. For authors, stabilizing chapter quality and adopting informative chapter titles can encourage continuous reading. Platforms, in turn, can boost retention by encouraging consistent contributions of "favorable" in-consumption comments—those providing scene-based insights, character evaluations, or pleas for new chapters.

Content Moderation with Shadowbanning (p. 1304)
Afrouz Hojati, Barrie R. Nault

Social media platforms face increasing pressure to moderate harmful content while preserving user engagement and free expression. We examine shadowbanning—a strategy that hides content without notifying the user—and compare it to traditional content removal. Our results show that if users only moderately believe that shadowbanning occurs, the platform benefits from a larger user base and higher profit, which also leads to greater social welfare than with content removal or no moderation. Shadowbanning allows the platform to reduce users' exposure to extreme content without deterring content creators, enabling more participation of users across the extremeness spectrum. However, outcomes depend on user beliefs and the accuracy of moderation technology. When users are highly suspicious of shadowbanning or when moderation tools are significantly imperfect, the platform's incentives—and the societal benefits—decline. These findings offer practical insights for platform designers and regulators: shadowbanning can be effective, but its benefits hinge on how transparently and accurately it is implemented. Policymakers should account for user perceptions and technological capabilities when evaluating or regulating opaque moderation strategies.

Attention or Sentiment: How Social Media React to ESG? (p. 1323)
Xiaoquan (Michael) Zhang, DaPeng Xu, Hong Hong, Kalok Chan

Despite the increasing attention associated with ESG investing, a natural yet overlooked question is how investors would react to the ESG performance of firms, given that investors' activities can ultimately impact firm performance. Utilizing data from multiple sources, we empirically find that ESG performance positively predicts future social media attention, but it has no predictive ability for future social media sentiment. Furthermore, the significant positive association between ESG and social media attention holds for both ESG downgrade events and ESG upgrade events. Moreover, the positive association between ESG and social media attention is driven by the environmental and social factors, whereas the uncorrelation between ESG and social media sentiment is determined by the social and governance factors. These findings urge managers to pay close attention to keeping and improving their firms' ESG advantage. They help managers shed light on the picture of the return of their firms' investment in ESG by taking investor behavior into account. Besides, ESG's inability in predicting social media sentiment also indicates retail investors' limited awareness of the importance of ESG currently. Therefore, it is pressing for regulators to take actions to boost the function of sustainable investing and ESG values.

When Gig Workers No Longer Gig: The Impact of California Assembly Bill 5 on the Online Labor Market (p. 1337)
Xunyi Wang, Yu-Wei Lin, Wencui Han, Liangfei Qiu

The rapid growth of the gig economy has intensified debates over how to classify workers. California's Assembly Bill 5 (AB5) reclassifies many independent contractors as employees, yet its practical effects on gig workers remain unclear. Using data from a major online labor platform and a difference-in-differences design, we provide one of the first empirical assessments of AB5's impact on workers' earnings. We find that California gig workers' monthly earnings rose relative to workers in other states, but this increase was driven by longer working hours, offsetting a drop in hourly pay. These results highlight the complex trade-offs of reclassification: although workers may gain access to employee benefits, they may also experience reduced hourly rates and increased workloads. For policymakers, our findings offer evidence to guide similar initiatives under consideration in other states and underscore the need for nuanced, context-sensitive regulation. For platform managers, the results suggest opportunities to align platform design and worker education with evolving labor laws to support compliance and sustainable growth. Together, these insights inform future policy and platform strategies aimed at balancing worker protection with flexibility in the gig economy.