Information Systems Research

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

Research Spotlights

To cite this article:
(2024) Research Spotlights. Information Systems Research 35(3):iii-x. https://doi.org/10.1287/isre.2024.resspot.v35.n3

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

INFORMATION SYSTEMS RESEARCH
Vol. 35, No. 3, September 2024, pp. iii–x
ISSN 1047-7047 (print), ISSN 1526-5536 (online)

Research Spotlights

https://doi.org/10.1287/isre.2024.resspot.v35.n3

More Than a Bot? The Impact of Disclosing Human Involvement on Customer Interactions with Hybrid Service Agents (p. 936)
Ulrich Gnewuch, Stefan Morana, Oliver Hinz, Ralf Kellner, Alexander Maedche

To leverage the complementary strengths of humans and artificial intelligence (AI) in online service encounters, firms have begun to use *hybrid service agents*: combinations of AI agents (e.g., chatbots) and human agents (e.g., service employees) behind a single interface. However, it is unclear whether firms should be transparent about behind-the-scenes employees working in tandem with an AI-based chatbot to serve customers. Against this backdrop, we investigated the impact of human involvement disclosure on customer interactions with hybrid service agents. Our findings suggest that disclosing human involvement before or during an interaction with the hybrid service agent leads customers to adopt a more human-oriented communication style. This effect is driven by impression management concerns that are activated when customers become aware of humans working in tandem with the chatbot. The more human-oriented communication style ultimately increases employee workload because fewer customer requests can be handled automatically by the chatbot and must be delegated to a human. These findings provide novel insights into how and why disclosing human involvement affects customer communication behavior, reveal its negative consequences for employees working in tandem with a chatbot, and highlight the potential costs and benefits of providing transparency in customer–hybrid service agent interactions.

A Design Theory for Transparency of Information Privacy Practices (p. 956)
Tobias Dehling, Ali Sunyaev

Public policy initiatives demand transparency of information privacy practices (TIPP) since the 1970s, at least, when a first version of the fair information practice principles (FIPP) was published. Decades later, information system (IS) providers are still struggling to meet this demand. Apparently, there is a tension between socially desirable guardrails for privacy practices (e.g., allowing for anonymity) and economic value creation (e.g., exploiting personalization surplus). To inform IS designs that can change such avaricious states to the benefit of consumers, we present a theoretical foundation (TIPP theory) explaining and prescribing what needs to be built to establish TIPP in the real world. On the one hand, TIPP theory explains why IS providers, who have privacy practices to hide, should not go beyond IS designs that meet legal demands. More importantly, TIPP theory explains and prescribes what IS providers need to build to actually establish TIPP and be able to differentiate from competitors by doing so. TIPP theory will be useful for all involved parties: IS providers are enabled to avoid privacy notice fallacies, public policy makers can give more details in their related guidance and laws published, and consumers can, ultimately, interact with IS that actually establish TIPP.

The Attention Economy: Measuring the Value of Free Goods on the Internet (p. 978)
Erik Brynjolfsson, Seon Tae Kim, Joo Hee Oh

We develop a framework to measure the value of free goods and services available on the internet. The conventional method of measuring consumer surplus based on monetary expenditures is ineffective because these goods' prices are predominantly zero. As the saying goes, time is money, and thus, our method addresses this challenge by quantifying the value of the time that consumers devote to consuming these free goods. We apply our model to data on the allocation of individuals' time among internet, television, leisure, and work, and this enables us to estimate the consumer surplus of free goods. We find that the average incremental welfare gain from the internet for each year between 2002 and 2011 was about $38 billion per year in the United States, which is equivalent to approximately 0.29% of the annual GDP. In contrast, if we had not considered the value of time, then the estimated annual incremental welfare gain would have been only $2.7 billion, which is barely 7% of the estimate derived from our model, which considers the time spent on consuming these goods. Our approach can be readily extended to the valuation of other zero-priced goods and services, such as television. In addition, our results show the importance of not only quantity but also quality (e.g., internet speed) in determining the welfare contributions of free goods.

Is a College Education Still Enough? The IT-Labor Relationship with Education Level, Task Routineness, and Artificial Intelligence (p. 992)
Dawei (David) Zhang, Gang Peng, Yuliang Yao, Tyson R. Browning

This study investigates the relationship between information technology (IT) and human labor through the lens of education level, task routineness, and artificial intelligence (AI). It makes use of a comprehensive data set covering 60 U.S. industries from 1998 to 2013. The results show that a college degree may not be sufficient, especially in routine-intensive industries, where a graduate degree may be increasingly necessary to meet the demands of computerized work environments. They also suggest that continuous education and IT skill development are essential for workers to thrive in a technology-driven economy. Governments should continue to support higher education to meet the demands driven by IT implementation, particularly as advanced technologies like AI emerge. Long-term training programs should be established to equip less-educated workers with the necessary skills to navigate technological changes. Additionally, policy makers should be cautious about the potential impact of advanced IT, such as AI, on employment rates and consider the consequences of increasing minimum wage levels, as it may incentivize firms to invest more in automation. Overall, this study emphasizes the need for a comprehensive understanding of the IT-labor relationship, guiding decision making for individuals, firms, educators, and policy makers in adapting to the challenges and opportunities presented by advancing technology.

Coordination in Multibrand, Multimedia Advertising: Is It Always a Good Thing? (p. 1011)
Wangsheng Zhu, Subodha Kumar, Vijay Mookerjee

The growing online retail market has led to the prevalence of multichannel retailing. Meanwhile, retailers are increasingly combining multichannel retailing with a multibranding strategy. Although this can further increase the retailer's sales, it brings new advertising challenges. Multibrand, multichannel retailers usually launch advertising campaigns for different brands on multiple media. Thus, their advertising efforts fall into a set of brand-media units. Each unit's advertising can affect the sales of all brands on all channels. To maximize the advertising returns, a retailer needs to coordinate the advertising expenditures on different units. We develop a stochastic differential equation model that helps the retailer estimate the impact of multimedia advertising on the sales of different brands on different channels. Afterward, we formulate the advertising optimization problems under four coordination strategies: noncoordination, brand coordination, media coordination, and global coordination. By solving the problem for each strategy, the retailers can obtain the optimal expenditure for each unit under that strategy. Finally, we compare the retailer's profits under four strategies. Interestingly, we find that brand or media coordination may result in a profit lower than noncoordination. Our findings provide insights regarding the selection of coordination strategies for multibrand, multichannel retailers with multimedia advertising campaigns.

Contextual Targeting in mHealth Apps: Harnessing Weather Information and Message Framing to Increase Physical Activity (p. 1034)
Nakyung Kyung, Jason Chan, Sanghee Lim, Byungtae Lee

Mobile technologies provide a unique opportunity for practitioners to identify users' real-time context and provide personalized interventions to influence their behaviors. However, less is known about a way to improve the effectiveness of mobile health intervention by using context information. This study provides design guidelines on how to use weather information with messaging formats to spur exercise. Through a field experiment that each participant experience different weather conditions in two different treatment periods under the gain or loss interventions, we found that the effects of gain or loss interventions under different weather conditions are heterogeneous. Loss intervention leads to higher fulfillment of exercise goals than gain intervention in sunny weather, whereas gain interventions are more effective than loss interventions in cloudy weather. In addition, we found that weather-based intervention can be used repeatedly over time without losing its effectiveness. Furthermore, we reveal that weather-based intervention is effective toward at-risk populations such as inactive individuals or lower income groups, serving as an mhealth solution that closes the health gap between the haves and have nots. Our findings provide useful guidelines for health service providers and health policymakers regarding how to effectively leverage contextual cues into mobile health intervention.

Does Help Help? An Empirical Analysis of Social Desirability Bias in Ratings (p. 1052)
Jinyang Zheng, Guopeng Yin, Yong Tan, Jianing Ding

Review-in-review (RIR) is a feature that allows viewers to generate positive or negative evaluations for primary quality evaluations of a product (e.g., ratings and reviews). This study reveals that it can cause social desirability bias in primary ratings: Reviewers who desire social recognition are driven to adjust their ratings (about 7.4% likelihood) to elicit more helpful responses and avoid unhelpful ones. This bias can be shown as distorted conformity to the prior rating

distribution or extremity, depending on the RIR types. The model identifies how bias magnitude correlates with users' social characteristics, thereby identifying vulnerable individuals. Platforms can incentivize less vulnerable users and remind susceptible ones to decrease the bias and can supplement rating conditional on the identified vulnerability extent (e.g., the distribution by the "independent" raters) to mitigate the bias's impact on rating viewers. The simulation analysis compares the bias under different counterfactual RIR system designs, finding a composite RIR system (e.g., helpful and unhelpful RIRs) partially neutralizes the bias, obviating the need to remove all RIR features. The model further adapts to evaluate underexplored RIRs forms and can provide a "de-biased" metric while preserving individual ratings.

Online Food Delivery Platforms and Female Labor Force Participation (p. 1074)
Jialu Liu, Siqi Pei, Xiaoquan (Michael) Zhang

Our research examines the impact of online food delivery platforms on female employment in South Korea, highlighting the often-overlooked positive externalities that digital platforms generate on the labor market and economics. We find that these platforms lead to an immediate and lasting increase in the female employment rate, with a notable increase in the female employment rate by 6.5%. The economic benefits derived from this rise in female employment account for 0.27% to the country's GDP, or 17 times the revenue of the online food delivery platform. Our analysis shows that such digital platforms offer a pathway for women to break free from traditional household roles, thus granting them more time and the opportunity to decide whether to join the labor market or stay at home. The freedom to be able to join the labor force promotes gender equality and fosters economic engagement. Policymakers need to recognize such indirect effects of technology-driven business models because they shape labor markets and drive economic outcomes. This research has vital implications for practice and policy, offering fresh perspectives on fostering female labor force participation and boosting economic growth. Empowering individual households and granting women greater agency can have profound effects, supplementing organizational measures. The study highlights the underestimated economic and societal value created by innovative technology platforms, emphasizing the need for policymakers to consider these spillover effects when assessing and regulating such platforms. Leveraging digital platforms can lead countries toward a more inclusive and prosperous future, dismantling barriers and promoting gender equality.

Healthcare Across Boundaries: Urban-Rural Differences in the Consequences of Telehealth Adoption (p. 1092)
Meizi Zhou, Xuelin Li, Gordon Burtch

We study the impacts of telehealth adoption on geographic competition among urban and rural healthcare providers. We consider a quasinatural experiment: states' entry into the Interstate Medical Licensure Compact, wherein the entry events facilitate healthcare providers to adopt telehealth technology. By analyzing a representative sample of providers, we first establish the Compact entry shock's validity and its positive effect on the supply of medical services. We then report evidence that there are service and payment shifts from rural providers to urban providers (i.e., urban providers are more likely to benefit from the Compact entry financially). Relying on patients' telehealth reimbursement claim data, we observe two mechanisms contributing to the revenue redistribution: the substitution and gateway effects of telehealth. Finally, we show that telehealth readiness and service quality moderate the impact of telehealth adoption. These findings speak to both potentially positive and negative consequences for welfare.

Consequences of Information Feed Integration on User Engagement and Contribution: A Natural Experiment in an Online Knowledge-Sharing Community (p. 1114)
Zike Cao, Yingpeng Zhu, Gen Li, Liangfei Qiu

This paper investigates the ramifications of information feed integration on user engagements and contributions in online content-sharing platforms by exploiting a natural experiment occurred in a leading knowledge-sharing platform that integrated informal social posts with professional knowledge content in one feed. Our results show that the juxtaposition of incongruous types of content increased mindset switching and cognitive strain, thus hurting user engagements. We also reveal a novel crowding-out effect, viz., the integration heightened concerns that posting informal social posts would dilute the contributor's professional image, thus inhibiting user contributions. Our findings hold important practical implications for all platforms that host (or are considering hosting) diverse types of user-generated content (UGC). Additional content curation tools can potentially enhance user engagement and retention, but their effectiveness hinges on a foundational and crucial element—the presentation format of heterogeneous content types. Essentially, the value of curating informal social posts in a knowledge-sharing platform would diminish when those content intrudes upon and conflict with the professional domain. This insight underscores that any UGC platforms, when adopting a diversity-oriented strategy, should pay close attention to heterogeneity between different content types for the purpose of optimizing user experiences and promoting user contributions.

Coopetition has been a common practice, especially among emerging markets. The coopetition relationship between a ride-sharing platform and a car-rental firm is distinct in that they operate under two different business models. Although the platform controls both its demand and supply by setting rider prices and driver wages, the car-rental firm operates under the traditional model with a fixed supply and cost structure. Both the platform and car-rental firm compete for riders seeking transportation. If the two cooperate, a driver is allowed to rent from the rental firm and drive for the platform; otherwise, only those owning personal vehicles are allowed to drive for the platform. We show that such supply-side cooperation intensifies demand-side price competition and decreases total revenue. Therefore, coopetition is mutually beneficial only when it leads to a significant decrease in supply costs. We find that the two firms are likely to form a coopetition relationship when the total rider market size is not high, the degree of rider substitutability between the two firms is low, and the platform has a significant market-size advantage over the rental firm. Coopetition between the platform and the rental firm benefits riders and hurts drivers, but it benefits society overall.

With the growing popularity of e-commerce, nearly every prominent retailer is aiming to turn omni-channel. One crucial decision in this pursuit is the identification of the joint assortment. In this study, we contribute by examining joint assortment and product prices for a retailer that sells products through both brick-and-mortar and online channels. Our analysis indicates that the optimal assortment should be thought of as a portfolio of two types of products: customized and omni-channel. Customized products are priced in such a way that they are targeted toward customers who prefer to shop from the channel the products are sold through. In contrast, omni-channel products are priced attractively so that all customers consider buying them. The relative mix of these products depends on how flexible customers are in shopping from the channel they do not prefer and the number of customers who prefer each channel. Additionally, we investigate whether the conventional wisdom of selling niche products through the online channel is always optimal. We find that this suggestion may be sub-optimal when the online channel has greater cost of including a product in the assortment and fewer preferring customers compared with the brick-and-mortar channel.

Many organizations have implemented internet monitoring to curb employees' non-work-related internet activities during work hours, commonly referred to as "cyberloafing." For managers, two primary considerations emerge: (1) the actual effectiveness of internet monitoring in diminishing cyberloafing and (2) any unintended side effects this monitoring might have on overall employee behavior. From a longitudinal field quasi experiment, we observed that although internet monitoring notably reduced cyberloafing because of amplified employee concerns about potential sanctions and privacy breaches, it unintentionally suppressed their organizational citizenship behavior (OCB). Moreover, a follow-up observation four months after introducing internet monitoring revealed that its capability to mitigate cyberloafing had weakened, yet the dampening effect on OCB continued. We conclude this paper by underlining the value of using internet monitoring as a feedback mechanism on employees' online behavior, rather than solely as a deterrence measure.

To reduce the availability of hacking tools for use in cybersecurity offenses, many countries have enacted computer misuse acts (CMA) that criminalize the production, distribution, and possession of such tools with criminal intent. Nevertheless, our research illuminates an unintended consequence: the chilling effect of CMA enforcement on legitimate cybersecurity discussions, some of which may be desirable for cybersecurity research, within online hack forums. More importantly, this study uniquely examines the chilling effect stemming from users' fear of legal harm. Drawing on decision-making theories related to choice under uncertainty, we derive new insights into how legal enforcement can suppress lawful acts and reveal the dynamics of social categorization online. Our research offers valuable insights for policymakers and forum administrators. Policymakers can use our findings to mitigate unnecessary uncertainty in legal enforcement such as CMA. This includes developing legal cases to prevent false prosecutions, implementing tailored communication strategies for inexperienced individuals, and considering supplementary measures like licensing and community recognition. A transparent

mechanism involving a neutral panel can also be established to ensure legal interpretations align with community norms. Forum administrators, on the other hand, can provide additional information and guidelines, foster responsible online environments, and align resources with professional standards to navigate the uncertain legal landscape and mitigate the chilling effect on knowledge-sharing.

Dinner at Your Doorstep: Service Innovation via the Gig Economy on Food Delivery Platforms (p. 1216)
Geng Sun, Yeongin Kim, Yinliang (Ricky) Tan, Geoffrey G. Parker

Despite the rapid growth of the online food delivery (OFD) market, the impact of its three-sided nature—encompassing consumers, restaurants, and gig drivers—on incentives and payoffs remains unclear compared to the traditional two-sided model. This study examines how OFD platforms make optimal choices in a competitive environment involving pricing and service quality. The analysis reveals that insights from two-sided platforms don't seamlessly translate to OFD markets. The triad nature of OFD can either dampen or heighten price competition in the buyer-seller market, altering subsidization dynamics for platforms. While conventional platforms suffer from negative network effects due to participation pressure, OFD platforms can adapt service strategies to mitigate this. However, introducing gig labor might not always benefit OFD platforms as it could trigger a prisoner's dilemma situation by empowering competing platforms. The study underscores the dependence of platform strategies on network effects' strength. As the gig economy rises, the employment status of gig workers garners controversy. The study demonstrates that implementing minimum wage regulations, while benefiting gig drivers, might diminish societal welfare. These findings offer guidance to policy makers aiming to balance gig workers' interests with overall societal concerns.

Consequences of China's 2018 Online Lending Regulation and the Promise of PolicyTech (p. 1235)
Yidi Liu, Xin Li, Zhiqiang (Eric) Zheng

Swift and unexpected shifts of financial regulations can have profound implications for the general population. This is evidenced by China's abrupt transition in its stance on P2P lending in 2018. Initially embracing these platforms, the abrupt regulatory pivot to widespread shutdowns. Our empirical research, drawing upon credit application data, demonstrates how this indiscriminate approach hindered economic development opportunities for a significant portion of borrowers, particularly the underprivileged. As a remedy, we advocate for the implementation of AI-driven regulatory frameworks. Rather than a monolithic approach to all borrowers, AI helps distinguish between real financial risks and those that can be managed. This nuanced strategy safeguards individuals' economic progression, while efficiently mitigating financial hazards. For policymakers and industry stakeholders, our findings underscore the importance of contemplating the broader ramifications of regulatory decisions and harnessing innovative methodologies, such as AI, to strike an optimal balance.

When Variety Seeking Meets Unexpectedness: Incorporating Variety-Seeking Behaviors into Design of Unexpected Recommender Systems (p. 1257)
Pan Li, Alexander Tuzhilin

In this paper, we study the consumers' variety-seeking behavior in recommender system applications and propose a comprehensive framework to measure such behavior based on past consumption records. The effectiveness of the proposed framework is validated through user questionnaire studies conducted at Alibaba, where our constructed variety-seeking measures match well with consumers' self-reported levels of their variety-seeking behaviors. We subsequently present a recommendation framework that combines the identified variety-seeking levels with unexpected recommender systems in the data mining literature to address consumers' heterogenous desire for product variety, in which we provide more unexpected product recommendations to variety-seeking consumers and vice versa. Through offline experiments on three different recommendation scenarios and a large-scale online controlled experiment at a major video-streaming platform, we demonstrate that those models following our recommendation framework significantly increase various business performance metrics and generate tangible economic impact for the company. Our findings lead to important managerial implications to better understand consumers' variety-seeking behaviors and design recommender systems. As a result, the best performing model in our proposed frameworks is deployed by the company to serve all consumers on the video-streaming platform.

Racial Discrimination and Anti-discrimination: The COVID-19 Pandemic's Impact on Chinese Restaurants in North America (p. 1274)
Chuang Tang, Shaobo (Kevin) Li, Yi Ding, Ram D. Gopal, Guanglei Zhang

The COVID-19 pandemic has seen a rise in racial discrimination against Asian communities, notably the Chinese population. Despite growing research on various aspects of the pandemic, there is a notable gap in understanding its behavioral impact regarding racial discrimination. This study delves into the manifestations of COVID-19–related racial discrimination and antidiscrimination efforts on online platforms using large-scale data sets from Yelp.com and SafeGraph. We specifically examined how the pandemic affected Chinese restaurants compared

with non-Chinese ones at different pandemic phases. Our findings are significant; the pandemic triggered an immediate surge in racial discrimination, resulting in a substantial decrease in customers visiting Chinese restaurants. Importantly, we applied advanced text mining and machine learning techniques to analyze user behavior, consistently revealing that increased discrimination prompted users to take antidiscrimination actions on online platforms. This research highlights a tangible form of racial discrimination through reduced patronage of Chinese restaurants and underscores the capacity of consumers to combat discrimination on online platforms. It calls for targeted policy interventions to address and prevent racial discrimination, particularly in the context of public health crises.

Development Trajectory of Blockchain Platforms: The Role of Multirole (p. 1296)
Tianyi Li, Xiaoquan (Michael) Zhang

Understanding the development trajectory of digital platforms is central to digital platform management. We develop a parametric model that investigates the development trajectories of blockchain platforms, accounting for the feedback between blockchains' utility change and people's adoption and abandonment behavior. We consider a typical blockchain participant to simultaneously play three roles on the platform, user, investor, and laborer, each contributing to blockchains' multi-faceted utility: providing service for transaction/interaction, providing a medium for digital investment, and providing workspace for online labor. The model describes a three-phase development trajectory for blockchain platforms: a chaotic initial stage, a rapid growth stage, and a mature stage of stable market cycles. The model was used to match 112 token price series, demonstrating robust performance across different fitting setups and outperforming existing models. The study identifies two temporal parameters, the time delay in quitting the platform and the holding time of the platform's token, that significantly differentiate blockchains' development trajectories. We extend the model to study forking events; results suggest that fork launch time is more important than forking amplitude in influencing the main chain's subsequent development and that forking can increase the exposure of the forked platform.

How Does Online Information Influence Offline Transactions? Insights from Digital Real Estate Platforms (p. 1324)
Zhengrui Jiang, Arun Rai, Hua Sun, Cheng Nie, Yuheng Hu

This study highlights the critical function that digital real estate platforms, like Zillow, serve in facilitating effective property transactions. They do this by transmitting vital property information from sellers to buyers, thereby enriching the value of offline deals. Our findings indicate that Zillow, as a source of information, is incredibly valuable for properties that deviate significantly from their neighborhood's average value, either above or below. It's particularly useful in conveying experiential details through images and textual descriptions. For potential buyers, Zillow is a trustworthy source of property information for estimating property value, especially when alternative sources of information are limited. This study underscores the necessity for sellers and their agents to effectively represent property information online, considering its significant impact on sale prices. This is especially true for unique properties and properties with notable experiential elements. Furthermore, our study suggests that real estate professionals need to modify their business practices to take full advantage of digital platforms and provide superior services to their clients. Finally, digital real estate platforms can use these insights to enhance their platform design by focusing on the collection and display of significant information, ultimately increasing the value provided to both buyers and sellers of properties.

Join Up or Stay Away? Coalition Formation for Critical IT Infrastructure (p. 1344)
Hong Guo, Yipeng Liu, Barrie R. Nault

We consider the formation of a coalition when districts invest in critical IT infrastructure that, if disrupted, can cause significant damage to security, the economy, public health, or safety. The benefits from these investments can spill over to other districts. Districts choose whether to participate in a coalition, and the coalition subsequently makes IT infrastructure investment decisions for those districts that join the coalition. These inside districts have superior interoperability in their spillovers relative to outside districts. We find that inside districts' resource levels decrease with the size of the coalition, and this size depends on the coalition's economies of scale and relative interoperability. Depending on these factors, any size coalition can be an equilibrium or socially optimal. In most cases, the socially optimal coalition size is larger than the equilibrium coalition. A subsidy or tax can incentivize the equilibrium coalition size and district investment levels to be socially optimal, providing a general solution to the provisioning of critical IT infrastructure. We use the European Union's Digital COVID Certificate program providing vaccine status information and the U.S. Government's Direct Project that supports the establishment of nationwide health information exchanges to illustrate elements of our model.

Uncertainty Reduction vs. Reciprocity: Understanding the Effect of a Platform-Initiated Reviewer Incentive Program on Regular Ratings (p. 1363)
Jingchuan Pu, Young Kwark, Sang Pil Han, Qiang Ye, Bin Gu

Many online platforms are now offering free samples to seasoned reviewers, hoping to get feedback. While these reviewers are given free samples to review, they also buy and review products themselves. The regular ratings for the purchased products are the majority. This brings up the question: Does receiving free products make them rate their personal purchases more positively? And if so, why? We explored two possibilities. First, *uncertainty reduction mechanism*: The idea that trying free samples makes buyers more confident in their purchases, leading to greater satisfaction and higher ratings for the purchased products; Second, *reciprocity mechanism*: The idea that reviewers might feel obliged to give better ratings as a "thank you" for the free samples or with the expectations of getting more free samples, which could introduce bias. Our research indicates that giving free samples mainly helps in reducing purchase uncertainty, making customers genuinely happier with their subsequent purchases. So, online platforms can benefit from this strategy, as it seems to uplift genuine positive reviews rather than create biased ones. However, it is still essential to monitor for any undue bias to maintain trustworthiness in reviews.

Dynamic Bayesian Network–Based Product Recommendation Considering Consumers' Multistage Shopping Journeys: A Marketing Funnel Perspective (p. 1382)
Qiang Wei, Yao Mu, Xunhua Guo, Weijie Jiang, Guoqing Chen

Recommender systems are widely used by platforms/merchants to find the products that are likely to interest consumers. However, existing dynamic methods still face challenges with regard to diverse behaviors, variability in interest shifts, and the identification of psychological dynamics. Premised on the marketing funnel perspective to analyze consumer shopping journeys, this study proposes a novel and effective machine learning approach for product recommendation, namely, multi-stage dynamic Bayesian network (MS-DBN), which models the generative processes of consumers' interactive behaviors with products in light of their stage transitions and interest shifts. In this way, consumers' *stage-interest-behavior* dynamics can be learnt, especially the variability in interest shifts. This provides managerial implications for practice. MS-DBN demonstrates significant performance advantage with general applicability by extracting the generalizable regularity during shopping journeys, which compensates the diversity and sparsity frequently observed in consumer behaviors. In addition, aided by the identification strategies integrated into the learning process, the latent variables in the model can be detected such that consumers' invisible psychological stages and interests in products can be identified from their observed behaviors, shedding light on the targeted marketing of platforms/merchants and thus enriching the practical value of the approach.

Retargeted vs. Generic Product Recommendations: When is it Valuable to Present Retargeted Recommendations? (p. 1403)
Xiang (Shawn) Wan, Anuj Kumar, Xitong Li

Online platforms/retailers widely use collaborative filtering (CF)-based generic product recommendations to improve sales. These systems recommend products to a consumer based on the product co-views and co-purchases by other consumers on the website but do not leverage the consumer's browsing data. Based on a field study on a U.S. fashion apparel and home goods retailer's website, we show that informing generic CF recommendations to individual consumers' browsing history can generate substantial additional sales. Specifically, we show that it is optimal to offer generic CF recommendations to a consumer if the consumer has not carted a product and recommend products he or she has seen in the previous sessions (retargeted recommendations) if he or she has carted a product. Our simulation results show that such recommendations could result in a 3% increase in total sales compared with conventional generic CF recommendations. Online platforms/retailers with detailed consumer browsing data can implement such recommendations to achieve higher sales.

An Onto-Epistemological Analysis of Information Privacy Research (p. 1422)
Heng Xu, Nan Zhang

Privacy is one of the most pressing concerns in the continuously evolving landscape of information technology. Despite decades of vigorous and multifaceted exploration in the interdisciplinary field of information privacy, a consensual or unifying theory remains elusive. Moreover, the complexities of issues surrounding privacy are frequently labeled as "too big to understand" in the public press. At this critical juncture, it is beneficial to delve deeper into the foundational assumptions that privacy scholars have about privacy phenomena. In this commentary, we offer a fresh perspective by drawing on Dreyfus' influential exegesis of the Heideggerian onto-epistemological framework to reflect on these assumptions. The perspective we offer yields three integrative recommendations for future privacy research to open to new research directions. We illustrate how these new directions could not only grow future privacy research but

also facilitate the design of more effective privacy-protection measures in practice.

Countering State-Controlled Media Propaganda Through Labeling: Evidence from Facebook (p. 1435)
Patricia L. Moravec, Avinash Collis, Nicholas Wolczynski

In an era dominated by social media, users are regularly exposed to propaganda, including efforts by authoritarian countries to undermine trust in government and health officials during elections and throughout the COVID-19 pandemic. Facebook and Twitter have taken steps to address this issue by adding labels indicating state-controlled media from certain countries like Russia, China, and Iran. This article investigates the effectiveness of state-controlled media labels in countering propaganda on social media, with a focus on Facebook. The researchers conduct two controlled online experiments and analyze field data surrounding Facebook's policy change in June 2020. The results indicate that state-controlled media labels can be effective in reducing engagement. However, the efficacy of the labels depends on users actively noticing them and the sentiment toward the country indicated in the label. Labels for countries with negative public sentiment showed a significant decrease in engagement, while those for positively perceived countries did not have the same impact. The study suggests that social media platforms should inform users about labeling policies and display labels prominently. Although propaganda will likely remain on social media, efforts to reduce its spread can be effective with proper implementation and awareness.

The Performative Production of Trace Data in Knowledge Work (p. 1448)
Aleksi Aaltonen, Marta Stelmaszak

Firms increasingly harness data that are created as by-products of information systems usage to evaluate and manage employees. However, such "trace data" can be a double-edged sword. The data can provide a whole new visibility into work practices but also, make work less transparent if the employees start to change their behavior to shape the data. We study this dilemma in the context of knowledge work that has traditionally eluded behavioral measurement. We show that when knowledge workers become aware of data collection and have an interest in how their work may be represented by the data, they start to actively perform the data. We identify different patterns by which employees shape work practices to produce trace data. The changes affect not only the actions and data of the focal employee but also, the actions and data of their colleagues and subordinates. Therefore, to fully realize the potential of trace data, managers may need to get involved in designing the data and to set a trace data policy that states how the data will be used in the organization.

Ontology-Based Intelligent Interface Personalization for Protection Against Phishing Attacks (p. 1463)
Fatemeh Mariam Zahedi, Yan Chen, Huimin Zhao

Millions of users on the Internet have fallen into phishing website traps. Detection tools are designed to warn users against such attacks, but often fail to achieve this purpose. One crucial reason behind this is that users rarely have a chance to interact and build a relationship with a detection tool that stealthily runs at the backend. A warning message on a rarely seen interface from such a tool hardly inspires users' trust in its authenticity and accuracy. In this study, we propose an ontology-based intelligent interface personalization (OBIIP) design for the warning interfaces of phishing website detection tools. We first constructed an ontology of warning interface elements (OWIE), which is a comprehensive knowledgebase for warning interface design. We then used OWIE in the design and creation of an OBIIP prototype and assessed it in a laboratory experiment and an online experiment. The results show the significant value of OBIIP in improving users' performance in terms of self-protection against website phishing attacks and building a stronger relationship with the detection tool in terms of trust in and use of the tool.

The Anchoring Effect, Algorithmic Fairness, and the Limits of Information Transparency for Emotion Artificial Intelligence (p. 1479)
Lauren Rhue

Emotion artificial intelligence (AI) is shown to vary systematically in its ability to accurately identify emotions, and this variation creates potential biases. In this paper, we conduct an experiment involving three commercially available emotion AI systems and a group of human labelers tasked with identifying emotions from two image data sets. The study focuses on the alignment between facial expressions and the emotion labels assigned by both the AI and humans. Importantly, human labelers are given the AI's scores and informed about its algorithmic fairness measures. This paper presents several key findings. First, the labelers' scores are affected by the emotion AI scores, consistent with the anchoring effect. Second, information transparency about the AI's fairness does not uniformly affect human labeling across different emotions. Moreover, information transparency can even increase human inconsistencies. Plus, significant inconsistencies in the scoring among different emotion AI models cast doubt on their reliability. Overall, the study highlights the limitations of individual decision making and information transparency regarding algorithmic fairness measures in addressing algorithmic fairness. These findings underscore the complexity of integrating emotion AI into practice and emphasize the need for careful policies on emotion AI.