Information Systems Research

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

Research Spotlights

To cite this article:
(2023) Research Spotlights. Information Systems Research 34(3):iii-ix. https://doi.org/10.1287/isre.2023.resspotv34n3

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

INFORMATION SYSTEMS RESEARCH
Vol. 34, No. 3, September 2023, pp. iii–ix
ISSN 1047-7047 (print), ISSN 1526-5536 (online)

Research Spotlights

https://doi.org/10.1287/10.1287/isre.2023.resspotv34n3

Personalized Ranking at a Mobile App Distribution Platform (p. 811)
Shengjun Mao, Sanjeev Dewan, Yi-Jen (Ian) Ho

Personalization is an emerging digital strategy to engage users across different business domains. It is defined as the capability to match content, products, and services to individual users based on the knowledge of their past behaviors and revealed preferences. It has shown its great potential across a variety of contexts, including search engines, recommender systems, targeted marketing, and more. In this study, we examine personalization on third-party mobile app platforms, which account for a $36 billion market in 2021. We develop a comprehensive structural framework for the personalized ranking of app impressions, leveraging revealed preferences embedded in consumer clickstream data. To improve platform revenues, the framework jointly accounts for consumer utility and cost per action margin. A series of policy experiments highlights the value of personalization to various extent. Remarkably, personalized rankings at the individual level outperform the current practice by 16.73%. This cost-efficient approach showcases how platforms can leverage routine consumer clickstream data to personalize the ranking of app impressions, thereby more effectively monetizing mobile app distribution.

Bystanders Join in Cyberbullying on Social Networking Sites: The Deindividuation and Moral Disengagement Perspectives (p. 828)
Tommy K. H. Chan, Christy M. K. Cheung, Izak Benbasat, Bo Xiao, Zach W. Y. Lee

Cyberbullying on social networking sites escalates when bystanders join in the bullying. Bystanders' joining-in behaviors reinforce the abuse, expose victims to a larger audience, and encourage further abuse by signaling their approval of the aggressive behavior. This study developed an integrative model that explains bystanders' joining-in cyberbullying behaviors on SNSs to offer actionable insights into reducing such harmful behaviors. We tested the model using 1,179 responses using a scenario survey study. Our findings suggest that IT artifacts (including digital profile, search and privacy, relational ties, and network transparency) activated two key mechanisms that lead to cyberbullying joining-in behaviors: (i) the deindividuation experiences that attenuate self-identity and put salience on group/social identity, and (ii) the moral disengagement practices that permit the exercise of cognitive maneuvers to justify group-interested choices that do not align with social standard. The findings explain why people who do not know each other gang up to bully a target on social media. Platform owners who wish to discourage bystanders from joining in undesirable activities may consider regulating how users could share and access digital resources in a social network and should acknowledge the influence of social identity in igniting, driving, and prolonging harmful online group behaviors.

Direct and Indirect Spillovers from Content Providers' Switching: Evidence from Online Livestreaming (p. 847)
Keran Zhao, Yingda Lu, Yuheng Hu, Yili Hong

Content providers in online social media platforms, particularly livestreaming, often switch content categories. We propose a theory-based framework to study the direct and indirect spillover effects of content switching for livestreamers—individuals who broadcast content through livestreaming platforms. Contrary to conventional wisdom, we propose two positive spillover effects that are unique to the social media platform setting: (a) the entrant streamers do not just increase competition among streamers, but they also bring their own viewers to the new category (direct spillover), and (b) the entrant streamers influence incumbent streamers' viewer size by boosting category visibility through indirect network effects (indirect spillover). We also propose that the two spillover effects are contingent on the size of the entrant streamers' follower base. Our findings show that average content switching is associated with a 1.3% net increase in direct net viewer flow and a 2.6% net increase in indirect net viewer flow. Our results provide managerial implications for livestreaming platforms on how to maintain a healthy and fair community among streamers, such as encouraging popular streamers to switch to a category with many promising small streamers.

Online health question-and-answer platforms (OHQPs), where patients post health-related questions, evaluate advice from multiple doctors and select their most preferred answer, are a prominent channel for patients to receive medical advice in China. They are gaining traction globally and have the potential to circumvent resource-based, geographic, or circumstantial barriers that limit access to care. At the same time, they are underregulated in contrast to brick-and-mortar points of care. Ours is the first study to evaluate the quality of healthcare advice promoted by these platforms and to provide insight into how patients respond to advice. We found that, though advice is generally good, patients cannot discern good advice from bad, choose poor advice (when offered) as often as good advice, and do so to a greater extent in vulnerable categories such as pediatrics, cancers/tumors, and internal medicine. Moreover, we found that OHQPs exacerbate *care avoidance*. Our findings suggest that platform owners and policymakers should ensure that *signals of expert consensus* are provided to patients to better assist their choices on OHQPs. We also unveiled bad actors on OHQPs, including drug promoters and spammers, indicating that stronger oversight and accountability mechanisms are needed. Physician peer reviewing and auditing can address both problems.

We are witnessing a quiet revolution—the rise of empowered users. These non-IT professionals increasingly seek to leverage the ever-expanding amount of organizational data for analytics to support their initiatives, decisions, and actions. All too often, however, the enthusiasm of these users collides against the harsh reality—many of them lack sophisticated IT skills, and they struggle to find/access relevant data, understand their meaning, and extract and adapt them to meet their needs. We propose a powerful way to support empowered users with a combination of conceptual models and narratives. Conceptual models are diagrams that accurately and succinctly represent rules and patterns captured in data. Although somewhat intuitive, these models alone do not suffice, as interpreting them still requires some specialized IT knowledge. Hence, we add narratives—stories written in natural language. The narratives intuitively explain some of the challenging aspects of the conceptual models. We conducted a series of experiments and interviews with empowered users to assess our idea. These studies show the value of the conceptual models with narratives for understanding organizational data. Our work unlocks an important missing puzzle for further user empowerment—a way for non-IT professionals to get the most out of the data.

This study examines whether and how an online service marketplace can leverage refund options endorsed by different parties (i.e., the platform or sellers) to address the "lemons" problem that is due to the intangibility, variability, and unreturnable nature of the services sought. We show that both platform refund insurance and a seller-guaranteed refund increase service demand, with platform refund insurance as the more effective option and hence having a more effective signaling mechanism, and that sellers with a better reputation or less popularity might benefit less from refund options. An investigation on further use of the more effective refund option, a "having platform refund insurance or being cast out" policy (i.e., retaining platform refund-insured sellers but expelling uninsured ones), reveals the effectiveness of this policy in filtering out low-quality sellers, shown as an improved quality of sellers on the platform due to new sellers' replacing those who were expelled, yet a cost (i.e., a loss in demand and consumer welfare) for the platform due to the changes in characteristics (e.g., price) of sellers. This cost, however, is lower than the benefit from the improved quality of the sellers, so that the platform's overall performance improves.

Blockbuster projects on crowdfunding platforms are those that have achieved outstanding and exceptional performance. Regarding the impact of blockbuster projects on crowdfunding platforms, there are two plausible yet opposing predictions: they may exhibit both a *negative* effect by monopolizing backer attention and resources and a *positive* effect by increasing the activeness of backer side. This tension could be further complicated, considering that blockbusters are inherently heterogeneous. Drawing on the cross-side network effects literature and integrating insights from the unique features of crowdfunding, we develop a theoretical framework highlighting the multidimensional view of blockbuster effects: blockbuster projects have an overall positive effect on the performance of concurrent projects (*overall effect*), and such positive effect is stronger for related blockbusters (*local effect*) and blockbusters emerging before the focal project (*temporal effect*). With data from a

leading crowdfunding platform, the empirical analyses largely support our proposed framework. Using off-platform data and backer data, we unveil the mechanism driving the blockbuster effects: blockbusters increase the popularity of the platform, whereby backers participating in the blockbusters tend to develop positive impressions and heightened expectations, so that they are likely to participate in other concurrent projects. Our findings have important implications for project creators and platform operators.

Positive Demand Spillover of Popular App Adoption: Implications for Platform Owners' Management of Complements (p. 961)
Mi Hyun Lee, Sang Pil Han, Sungho Park, Wonseok Oh

As platform owners interact with end users and complementors, their demand-side characteristics and performance affect the overall value creation of ecosystems. This research investigated how the emergence of popular complements on a mobile communication platform impacts the usage of other complements by the platform's end users and how platform owners can benefit from such demand spillovers. Using individual user-level app usage data, we empirically demonstrated how the presence of a popular app alters the demand structure of a platform through changes in the usage of other apps operating within it. The findings reveal that popular app adoption by users increases the app usage, excluding the usage of popular apps, only within the platform offering a popular app, supporting the existence of positive spillovers from popular complement adoption on a platform. Such positive within-platform spillovers occur through the increased usage of both existing and new apps, with stronger positive spillovers for new apps. These findings imply that platform owners can reap benefits by coordinating the launch of new complements and the promotion of less-known counterparts to end users with the emergence of a popular app. All these results shed light on how platform owners can manage their complements and create value from popular complements.

Pushing Yourself Harder: The Effects of Mobile Touch Modes on Users' Self-Regulation (p. 996)
Yang (Alison) Liu, Zhenhui (Jack) Jiang, Ben C. F. Choi

Mobile health interventions are widely used to facilitate individuals' management of their health behavior. A notable issue is that health interventions with obvious persuasive intent may cause negligence and reactance. In this study, we propose a subtle but powerful way to bolster self-regulation in maintaining healthy behavior by leveraging embodied interaction design. Our study shows that bodily actions in interacting with digital devices can trigger thoughts about prior associated experiences and, thus, be strategically designed to affect individuals' judgments, decisions, and behavior. Specifically, in three experiments, we find that firmly pressing a touchscreen during mobile interaction (as compared with gently tapping a touchscreen) can activate users' approach motivation and, thus, induce their preference for a healthy over a tasty beverage, lead to more challenging exercise goals and more exercise, and reduce personal hygiene lapses after receiving hygiene education. Hence, designers of digital health products may consider designing interaction with pressing gestures to facilitate users' self-regulation and attainment of health-related goals. Policymakers can also encourage the adoption of relevant app designs to improve citizens' health wellbeing.

Sequential IT Investment: Can the Risk of IT Implementation Failure Be Your Friend? (p. 1017)
Vidyanand Choudhary, Mingdi Xin, Zhe Zhang

Information technology (IT) investment often faces the risk of implementation failure. Such risk has important and intriguing impacts on firms' investment and competitive strategies. We investigate these effects in a context in which two firms, a leader and a follower, invest in a cost-reducing IT sequentially, and IT implementation may fail. We show that the risk of IT implementation failure benefits the follower and can benefit or hurt the leader. It impacts the firms' investment and profits through three distinct effects: mitigating competition, weakening the leader's first mover advantage, and creating differentiation opportunities. The follower may have information about the leader's IT investment before making its investment. The leader gains a first mover advantage when the follower makes its investment after knowing the leader's IT investment level or implementation outcome. So pioneering IT adopters should announce their investment early to influence followers. In some cases, the follower can benefit from knowing the leader's implementation outcome. The leader expects a higher profit than the follower, and early investment is beneficial unless implementation risk declines significantly with time. Finally, a sequential IT investment schedule in which the follower makes its investment after the leader's investment level and outcome are known produces the highest social surplus.

Information Systems Research for Smart Sustainable Mobility: A Framework and Call for Action (p. 1045)
Wolfgang Ketter, Karsten Schroer, Konstantina Valogianni

Transportation is a backbone of modern globalized societies. It also causes approximately one third of all European Union and U.S. greenhouse gas emissions, represents a major health hazard for global populations, and poses significant economic costs. However, rapid

innovation in vehicle technology, mobile connectivity, computing hardware, and artificial intelligence (AI)-powered information systems heralds a deep socio-technical transformation of the sector. The emergence of connected, autonomous, shared, and electric (CASE) vehicle technology has created a digital layer that resides on top of the traditional physical mobility system. This article contributes a framework to direct research and practice toward leveraging the opportunities afforded by CASE for a more efficient and less environmentally problematic mobility system. The authors propose seven overarching dimensions of action. These range from designing real-time digital coordination mechanisms for the management of mobility systems to developing AI-powered real-time decision support for mobility resource planning and operations. Per each dimension, concrete angles of attack are suggested which, we hope, will spur structured engagement from both researchers and practitioners in the field.

Reidentification Risk in Panel Data: Protecting for *k*-Anonymity (p. 1066)
Shaobo Li, Matthew J. Schneider, Yan Yu, Sachin Gupta

Market research companies collect extensive data on purchasing, travel, and app and media usage behaviors of consumers, prescriptions written by physicians, and so forth. Although the companies provide assurances of anonymity to the study participants, there is a significant concern about the vulnerability of these data. Could a motivated intruder match the pattern of purchases with the name and other personal and potentially sensitive details of an individual? We find that 17% to 94% of market research panelists in 15 frequently bought consumer goods categories are subject to high risk of reidentification through a potential record linkage attack based on their unique purchasing histories even when their identities are anonymized. We also demonstrate that the risk of reidentification in such data are vastly understated by the conventional measure, unicity, and propose a new measure, termed "sno-unicity." To protect the privacy of panelists, we consider the well-known privacy notion of *k*-anonymity and develop a new approach called "graph-based minimum movement *k*-anonymization" that is designed especially for retaining the usefulness of panel data. We show that our approach works well in protecting participants' privacy without substantially altering the information that marketers need for sound marketing decisions.

A Theory of Information Compression: When Judgments Are Costly (p. 1089)
Richard Thomas Watson, Kirk Plangger, Leyland Pitt, Amrit Tiwana

How useful to tourists are thousands of reviews of different five-star hotels in a city on a travel website when the mean rating is 4.5, and all the five-star hotels score around the mean? How insightful are reviews of physicians on a physician review website to potential patients when the ratings cluster tightly around an average for all physicians? Are there costs to the physicians, the patients, and to society as a whole? When all the students at a university score "A" grades on most courses, are there consequences for the university, the students, and potential employers? This paper calls the "clustering around a mean" phenomenon "information compression" and the systems in which it occurs (e.g., universities, students, employers) "judgment networks." When there is extensive information compression in a system, measures such as ratings or grades have little value for decision makers. When all five-star hotels in a city score an average of 4.5 does it really matter which one a traveler chooses? The paper introduces a way of measuring information compression. It also suggests ways for organizations to overcome the negative consequences of information compression for themselves and their various stakeholders.

Single-Sourcing vs. Multisourcing: An Empirical Analysis of Large Information Technology Outsourcing Arrangements (p. 1109)
Ravi Bapna, Alok Gupta, Gautam Ray, Shweta Singh

As the information technology (IT) services landscape matures, clients are increasingly adopting multisourcing arrangements that involve multiple vendors. Although a large body of IS literature addresses issues of whether to outsource (to a single vendor), what types of contracts to use, and how to achieve optimal relational governance, little is known about the antecedents and consequents of the single versus multisourcing decision. Moreover, although conceptual and analytical models of single-sourcing versus multisourcing have been developed, there is no empirical IS research using a large-scale data set with rigorous econometric analysis that examines the antecedents and consequents of multisourcing in the IT context. This paper fills this void, using the transaction cost economic lens and a data set of 49,057 large IT outsourcing arrangements that spans multiple industries and dates back 25 years. For managers who plan to multisource IT outsourcing arrangements, this research provides guidance to minimize exchange hazards through a better understanding of the relationship between sourcing choice, client IT outsourcing capabilities, the competitiveness of the vendor landscape, and the number of IT services in an IT outsourcing arrangement. We provide empirical evidence that the choice between single-sourcing and multisourcing is material to the performance of outsourcing contracts, as an incorrect sourcing choice is likely to result in negative contract outcomes.

Sales on the e-commerce platform in the United States have experienced explosive growth and are projected to surpass $740 billion in 2023. The expansion of the platform's traditional role as a reseller into an online marketplace and the introduction of its own brand products have stoked a huge fear among the incumbent sellers. The platform's unfair anti-competitive practice further aggravates the situation. Consequently, politicians and regulators have proposed prohibiting platforms from introducing own brand products to protect the incumbent sellers. This study addresses two questions of critical interest to the policymakers and the incumbent sellers. First, how does the platform's introducing its own brand product affect the incumbent sellers? Second, how effective is the proposed policy in terms of protecting the incumbent sellers? We examine the impact of the platform's own brand introduction on the incumbent sellers under two prevailing sell-on and sell-to pricing contracts. We find that the proposed legislation "that prohibits platforms from both offering a marketplace for commerce and participating in that marketplace" does not have the desired outcome of helping the incumbent sellers. Instead, it forces the platform to adopt only the sell-to contract with its own brand introduction which hurts the sellers under most market conditions.

Customers in sales processes increasingly encounter automated sales agents that complement or replace human sales agents. Yet, little is known about whether, how, and why customers respond to automated agents in contrast to human agents across successive decision stages of the same sales process. Even less is known about customer responses to combinations where both agents assume distinct roles and focus on complementary tasks that are traditionally performed by only one single agent. Against this backdrop, this paper explores the influence of increasingly common sales representative types on customer decisions across sales stages. Our findings demonstrate that customer responses to automated (versus human) sales agents are not stable in sales processes and instead, shift as customers move across sales stages. What is more, the paper shows that combinations of sales agents versus single sales agents do matter, yet their differential effects depend on contextual features of the sales setting. These insights are important because vendors may assume that a certain type of sales agent is always more appreciated by customers, whereas in fact, different sales agent types bring distinct attributes to the table, and customers' appreciation of these attributes shifts across sales stages.

Should hospitals source electronic health record (EHR) applications from a single vendor or multiple vendors to deliver high-value care? This is an on-going debate that has assumed greater importance with the emergence of value-based care initiatives in recent years. We contribute to this debate by studying hospitals' EHR sourcing strategies based on their integration with physician practices and its impact of the value of healthcare. Drawing on a nationwide sample of U.S. hospitals, our research shows that there is no one-size-fits-all solution. A single-sourced EHR can facilitate information sharing and subsequently improve healthcare value. However, this phenomenon is pronounced only when hospitals and physician practices are tightly integrated. Hospitals that adopt an EHR system from one vendor may not realize its full interoperability benefits without tighter practice integration because of incentive misalignment. Hence, it is important for healthcare providers to prudently choose their EHR sourcing strategies based on their practice integration model. We also provide a novel optimization-based methodology to measure healthcare value using data envelopment analysis. As the industry strives to achieve a vision of high-quality care at a lower cost, our research provides a roadmap for practitioners and policymakers to thrive in a value-based care environment.

As artificial intelligence (AI) solutions are being rapidly deployed, they increasingly compete with human labor. This study examines designers' strategies in response to the threat from the introduction of an AI system for simple logo designs in a crowdsourcing design platform. We find that, although designers with lower abilities are more likely to exit the platform, designers with higher abilities move away from the locus of threat in the lower-tier contests and switch to more-complex design contests after the introduction of the AI system. More interestingly, we find that, although unsuccessful designers respond to the threat from AI by increasing their participation across multiple contests, successful designers become more focused (i.e., they substantially increase the number of submissions within a contest) and more quality oriented (i.e., they increase emotional

content and complexity of their designs) after the AI launch. Our findings show how designers can learn from the behaviors of the more successful designers to differentiate themselves from AI systems by leveraging the more-abstract design attributes. Platform operators would benefit from adopting better segmentations strategies: with AI solutions for simple design tasks, hybrid AI + human solutions for less-complex design tasks, and skilled human designers competing primarily for the more-complex design tasks.

Decide Now or Later: Making Sense of Incoherence Across Online Reviews (p. 1211)
Dezhi Yin, Triparna de Vreede, Logan M. Steele, Gert-Jan de Vreede

Consumers read online reviews to decide whether to buy a product. Extensive research examines what makes a single review helpful or credible, yet there is very limited understanding of how a *collection* of reviews facilitates purchase decisions. Such understanding is critical because consumers rarely consult all reviews or a single review. They often start by reading the "top reviews" that a website highlights, then deciding whether to read additional reviews and how many. This paper investigates how inconsistency among top reviews affects a consumer's purchase deferral—the likelihood to decide immediately or defer the decision until after obtaining more information. We found that, if different reviewers disagree on their opinions about the *same* feature of a product, consumers are more likely to defer the purchase decision and consult more reviews. Further, this effect is weaker when reviewers provide specific details about their needs or use of the product along with their opinions. This work provides guidance to review platforms on how to select and present a set of top reviews. Our findings also inform retailers and product manufacturers on how to focus their attention in dealing with reviews and when a focus only on the top reviews is not sufficient.

Value of Information Sharing via Ride-Hailing Apps: An Empirical Analysis (p. 1228)
Kyung Sun (Melissa) Rhee, Jinyang Zheng, Youwei Wang, Yong Tan

This study examines the effects of an information sharing restriction policy that restricts taxi drivers' access to ride requests via ride-hailing apps. We show that the policy significantly decreases the ridership of an affected taxi fleet during times of enforcement but significantly increases the demands at some times of nonenforcement after launch. Furthermore, the traffic on public transportation, including metro, bus, ferry, and park & ride, and the congestion on the surface roads and expressways significantly increase after launching the policy during both enforcement and most nonenforcement times. We also show that the profitability of taxi fleet decreases after the restriction, which supports the notion that information sharing via ride-hailing apps enables them to match not only with more orders but also with those of higher marginal profit. These findings suggest that information sharing via ride-hailing apps can improve the utilization of existing taxi capacity, which further alleviates traffic during alternative times and the burden placed on alternative transportation modes. Policymakers and platform managers should dissect the value of information sharing from that of other aspects (e.g., changes in supply) in on-demand platforms and design policies that more specifically restrict the harmful aspects rather than restricting the use of such apps.

Adjusting Skillset Cohesion in Online Labor Markets: Reputation Gains and Opportunity Losses (p. 1245)
Marios Kokkodis

In online labor markets, contractors' ability to charge for their services largely depends on their skills. To keep up with shifting labor market needs, contractors often expand their skills with new skills. When the new skills are similar to the contractors' current skills, they often increase skillset cohesion (i.e., the average similarity of skills in a skillset). However, when the new skills have little similarity with contractors' current skills, skillset cohesion decreases. How do such adjustments of skillset cohesion affect contractor value in digital workplaces for short-term work? We argue that skillset adjustments affect market value through changes in the contractor's perceived reputation on the new skills and the additional job opportunities that new skills create. We hypothesize that compared with skills that decrease cohesion, skills that increase cohesion result in reputation gains and opportunity losses. Analysis of a panel data set of 47,638 tasks illustrates that for hourly wages, reputation gains are smaller than opportunity losses; hence, all else being equal, increasing skillset cohesion has a relatively negative effect on wages. However, for hiring rates, the opposite is true: Increasing skillset cohesion increases contractor hireability.

Complementarity Between Investment in Information Technology (IT) and IT Human Resources: Implications for Different Types of Firm Innovation (p. 1259)
Feng Guo, Yijun Li, Likoebe M. Maruping, Adi Masli

Innovation is an important means by which firms generate new revenue streams for topline growth. Given the information intensity of innovation, a key question that managers face is: How do we organize our IT-related resource investments to promote innovation performance? This is a consequential question because different types of innovation incur varying levels of upside potential, failure risk and investment costs. In this research, we aim to answer this question by distinguishing four types of innovations—incremental, radical,

non-IT related and IT related—that firms aim to produce and examining the extent to which each of these types benefits from complementary investments in IT human resources and IT compared with either type of investment alone. Using patent data from the U.S. Patent and Trademark Office, our analysis of 36,812 firm-year observations reveals that complementary investments in IT human resources and IT are associated with a 117% increase in incremental innovation, a 155% increase in radical innovation, a 173% increase in non–IT-related innovation and a 161% increase in IT-related innovation. The results of this study can be helpful in enabling managers to make informed decisions about how to organize and get the most out of their resource investments based on their innovation objectives.

The Attraction Effect in Crowdfunding (p. 1276)
Markus Weinmann, Abhay Nath Mishra, Lena Franziska Kaiser, Jan vom Brocke

Crowdfunding has reached enormous popularity among entrepreneurs, yet many struggle to fund their projects. For example, only 39.8% of all Kickstarter projects get funded. Why? Entrepreneurs need to set a predefined funding goal, which needs to be reached to get the project funded. If not reached, they do not get the money. To reach the funding goal, entrepreneurs can design reward menus from which backers can choose low-priced to high-priced rewards to support a given project. We show that simply inserting a so-called decoy option—an option in the first place that does not make sense but makes higher-priced options more attractive—can lead to significantly more backers choosing a high-priced option (up to 28% more high-priced choices). This makes it more likely for entrepreneurs to reach their predefined funding goals and thus start their designated projects.

Bots with Feelings: Should AI Agents Express Positive Emotion in Customer Service? (p. 1296)
Elizabeth Han, Dezhi Yin, Han Zhang

The rise of emotional intelligence technology and the recent debate about the possibility of a "sentient" artificial intelligence (AI) urge the need to study the role of emotion during people's interactions with AIs. In customer service, human employees are increasingly replaced by AI agents, such as chatbots, and often these AI agents are equipped with emotion-expressing capabilities to replicate the positive impact of human-expressed positive emotion. But is it indeed beneficial? This research explores how, when, and why an AI agent's expression of positive emotion affects customers' service evaluations. Through controlled experiments in which the subjects interacted with a service agent (AI or human) to resolve a hypothetical service issue, we provide answers to these questions. We show that AI-expressed positive emotion can influence customers affectively (by evoking customers' positive emotions) and cognitively (by violating customers' expectations) in opposite directions. Thus, positive emotion expressed by an AI agent (versus a human employee) is less effective in facilitating service evaluations. We further underscore that, depending on customers' expectations toward their relationship with a service agent, AI-expressed positive emotion may enhance or hurt service evaluations. Overall, our work provides useful guidance on how and when companies can best deploy emotion-expressing AI agents.