



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(3):v-xiv. <https://doi.org/10.1287/isre.2026.resspot.v37.n3>

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>

Research Spotlights

<https://doi.org/10.1287/isre.2026.resspot.v37.n3>

Copyright: © 2026 INFORMS

How to Tell a (News) Story? Quantifying the Impact of News Format and Storytelling on Engagement (p. 1363)

Debashish Ghose, Susan Mudambi, Subodha Kumar, Joydeep Srivastava

Social media has transformed how people consume news, making storytelling design as important as headline design. We examine three storytelling design features—narrativity (linear versus summary first), emotional sequence (good to bad versus bad to good), and reading level (simple versus complex)—and show that their effects on engagement depend on both news format and audience motivation. We find that effective engagement requires conditional rather than universal rules. For satirical content, humor enhances engagement mainly among motivated audiences, where higher narrativity and complex language paired with bad-to-good emotional sequences work best. For less motivated audiences, simpler satire is more effective when presented in summary-first form with good-to-bad sequences. For traditional news, prior work suggests that simple language helps; whereas we find this to be generally true, our results show that when complex language is unavoidable, pairing it with linear narrativity and bad-to-good sequences can enhance engagement. These results provide guidance for publishers who must balance clarity, complexity, and audience expectations. Beyond news, our findings generalize to videos, podcasts, and interactive media, where storytelling design similarly shaped engagement. We also demonstrate how large language models can generate controlled story variations, enabling creators to scale production and platforms to optimize recommendations in audience-specific ways.

Collaborative Intelligence in Sequential Experiments: A Human-in-the-Loop Framework for Drug Discovery (p. 1390)

Jinghai He, Cheng Hua, Yingfei Wang, Zeyu Zheng

Drug discovery is a complex process that involves sequentially screening and examining a vast array of molecules to identify those with the target properties. This process faces challenges because of the vast search space, the rarity of target molecules, and constraints imposed by limited data and experimental budgets. To

overcome these challenges, we propose a human-centered human–algorithm collaboration framework. Notably, both the algorithm and humans have substantial knowledge gaps. The algorithm proposes, and human experts retain decision rights to approve, revise, or override. Our design artifact leverages dual-process theory for attention management, surfaces meta-knowledge as a shared state for coordination, and optimizes a joint team objective. In real-world drug discovery tasks, the human–artificial intelligence (AI) team consistently outperforms all baselines, including human-only and AI-only methods. This demonstrates complementary performance. These findings illustrate that the optimal use of AI in complex decision-making environments is to augment, rather than replace, human expertise. Managers should balance AI integration with the cultivation of deep domain expertise, ensuring that technology enhances decisions without eroding uniquely human strengths. For practice and policy: preserve human agency with clear decision rights, make model uncertainty explicit, and use AI to route scarce expert attention to the most informative experiments.

To Claim or Not To Claim? Hidden Costs of Business Page Claiming (p. 1416)

Jong Youl Lee, Mikhail Lysyakov, Huaxia Rui

Many digital platforms encourage small business owners to “claim” their pages to improve visibility and connect with customers. Yet, despite its apparent benefits and zero financial cost, many pages remain unclaimed. Using a unique data set from Yelp and a staggered difference-in-differences design, this study reveals a hidden downside of business page claiming; average customer ratings drop by 10.3%, driven by more one-star reviews and fewer five-star reviews. Customers also write longer, more negative reviews and address owners more directly about service issues. These findings indicate that claiming a business page signals owner presence and raises customer expectations for responsiveness, which many small businesses may not be equipped to meet. For practitioners, the results highlight that business page claiming, although free, is not costless; it creates reputational risk if service quality or responsiveness falls short of heightened expectations. Owners should claim their pages only when ready to

actively monitor feedback or engage in managerial responses. Platforms should also communicate these potential consequences to businesses and design tools that help owners manage customer interactions more effectively.

Agent-Based Data Curation Practices: Customer Responses to Human vs. Algorithmic Data Requesters in Established Business-to-Business Relationships (p. 1434)

Martin Adam, Abhay Nath Mishra, Alexander Benlian

With the increasing value generated through data curation and the rise of artificial intelligence (AI) agents that act as human agents, vendor companies in established business-to-business relationships increasingly delegate data curation tasks to algorithmic data requesters (ADRs) instead of human data requesters (HDRs). Using a randomized field experiment with a European pharmaceutical company and a follow-up online experiment, we show how customers respond to ADRs versus HDRs across two tasks: data enrichment (i.e., adding new information) and data reconciliation (i.e., updating existing records). For data enrichment, customers are more likely to agree and complete requests sent by ADRs because they expect lower effort. For data reconciliation requests, customers lean toward HDRs, reflecting stronger accuracy concerns. Interestingly, we observe only a marginal advantage in completion rates for HDRs. These findings advise vendors to match agents to data curation tasks; they should deploy ADRs for data enrichment to reduce customer burden, and they should use HDRs for data reconciliation to address error concerns. Relatedly, vendors should craft emails that match the messaging context; they should frame data enrichment messages around convenience and benefits, and they should frame data reconciliation messages around accuracy, auditability, and risk reduction.

Working Daily, Paid Monthly? Effects of On-Demand Wage Access on the Financial Engagement of Low-Wage Workers (p. 1463)

Jihye Kim, Seokchae Yoon, Sunghun Chung, Wonseok Oh

Low-wage workers often face liquidity constraints, relying on costly short-term credit options such as payday loans. On-demand wage access (OWA) platforms offer a promising alternative by allowing workers to access earned wages before traditional pay cycles. This study examines whether OWA enhances financial engagement among low-wage workers and investigates the mechanisms driving these effects. Using transaction data from about 4,000 users of a U.S.-based OWA platform, we employ a difference-in-differences approach with staggered adoption to identify behavioral changes. To enrich our analysis, we supplement these data with evidence from an online experiment, a survey of OWA

users, and semistructured interviews exploring psychological and contextual mechanisms. Our findings show that OWA adoption increases monthly saving frequency by 3.7%, time spent monitoring financial dashboards by 12.9%, and goal-setting activity by 1.3%. These effects are amplified in regions with lower minimum wages or limited banking access but weakened among users who frequently incur fees for instant withdrawals. We identify *self-empowerment*—beyond mere self-efficacy—as the key mechanism enabling a shift from reactive to proactive financial management. By demonstrating how digital affordances foster behavioral change within structural constraints, this research offers actionable insights for employers, OWA providers, and policymakers seeking to promote financial inclusion.

Knowing (Not) to Know: Explainable Artificial Intelligence and Human Metacognition (p. 1485)

Moritz von Zahn, Lena Liebich, Ekaterina Jussupow, Oliver Hinz, Kevin Bauer

Many organizations seek to combine human expertise with explainable artificial intelligence (XAI), but they often overlook a core requirement for effective collaboration; humans must understand their own abilities. This understanding of one's own abilities is referred to as *metacognition*, which captures how well individuals monitor and regulate their own decision making. In two experiments with real estate and lending professionals, we find that XAI improves human metacognition by reducing their overconfidence. As a result, experts can better delegate decisions to an artificial intelligence (AI) and interact with it more effectively, improving the collaborative performance. These effects occur primarily when XAI highlights differences between human reasoning and AI logic. Our findings demonstrate that metacognition is a key mechanism through which XAI affects decision outcomes and offer guidance for organizations deploying (X)AI under growing transparency and accountability mandates, such as those in the European Union Artificial Intelligence Act.

The Effects of the FTC Policy and Affiliation Disclosures on Product Review Video Engagement: Evidence from YouTube (p. 1508)

Jingwen Zhang, Stephanie Lee, Yong Tan

Affiliate marketing on social media involves content creators posting product reviews with affiliate links, through which they earn a commission from resulting purchases. The Federal Trade Commission (FTC)'s disclosure guidelines require creators to reveal affiliate relationships alongside product reviews. This paper examines how the FTC policy affects viewer engagement with affiliated content, defined as product review videos containing affiliate links. We find that after the FTC policy implementation, viewer engagement with

affiliated content significantly decreases relative to non-affiliated content. However, affiliation disclosure moderates this effect: after the policy, affiliated content with disclosures receives higher engagement than affiliated content without disclosures. The mitigating effect of disclosures varies across content creator and video characteristics. The mitigating effect is more pronounced for more experienced creators, more popular creators, and more negative videos. This paper provides important empirical evidence on disclosure policies' effects on viewer engagement, addressing a notable gap in the literature. Our findings also have important implications for policymakers designing effective disclosure regulations, content creators managing compliance while maintaining viewer engagement, and social media platforms implementing disclosure features. This paper finds that whereas disclosure policies may lower engagement with affiliated content, proper disclosures can signal credibility and increase engagement.

A Deep Learning Approach for Predicting FDA's 510(k) Medical Device Recalls Using Device Citation Relationships (p. 1526)

Yi Zhu, Soumya Sen, Alexander Everhart,
Pinar Karaca-Mandic

More than 90% of medical devices in the United States are approved through the Food and Drug Administration's 510(k) pathway, primarily based on demonstrating the equivalence of new devices (known as applicant devices) to previously cleared devices (known as predicate devices). However, safety concerns are raised as applicant devices cleared this way may be more prone to recalls that relate to substantial patient harm and financial strain on the healthcare system. In response, this work introduces a data-driven information technology approach to predict medical device recalls, aiming to alleviate these safety concerns by augmenting human decision making. The approach primarily uses the characteristics of the network formed by predicate device citation relationships (predicate network). It uses deep learning to tackle three design challenges: learning the predicate network structure, capturing the temporal patterns of predicate network characteristics, and accounting for dependencies across the predicate citation history. Based on 45,398 medical devices cleared between 2003 and 2020, the approach substantially improves recall prediction accuracy and timeliness compared with existing state-of-the-art approaches. The improved recall-prediction performance and insights into performance variations across device categories provide opportunities to preemptively react to potential recalls and improve the safety of devices cleared through the 510(k) pathway.

Asymmetric Learning Effects of Chief Information Officer Outside Board Appointments: Cybersecurity Implications for Sender and Receiver Firms (p. 1550)

Justin Short, John D'Arcy, Yili Hong

Cybersecurity failures are increasingly costly, prompting companies to recruit chief information officers (CIOs) from other firms to their boards. This study examines whether there are any impacts on a firm's cybersecurity when (a) firms allow their own CIO to serve on an outside board; and (b) when firms appoint a CIO from another company to their own board. Using CIO-firm-year observations, we compare two pathways: (1) receiver firms that appoint an external CIO to their board, and (2) sender firms whose own CIO serves on another company's board. The findings show asymmetric effects. Receiver firms experience fewer data breaches, suggesting that external CIOs effectively transfer cybersecurity expertise and practices. In contrast, sender firms face higher breach risk, as CIOs who serve externally appear to prioritize educating the recipient firm over acquiring new insights for their home firm. This risk intensifies when the external firm lacks strong cybersecurity practices but is mitigated when the home firm has a dedicated chief information security officer. Conversely, receiver firms benefit most when the sending firm has strong cybersecurity capabilities—or even a past breach—because negative events create valuable lessons. The results offer actionable implications: firms should strategically recruit outside CIOs to improve board-level cyber capabilities and carefully weigh the risks before permitting their own CIOs to serve externally. Policymakers should consider mechanisms that incentivize effective cybersecurity knowledge transfer across board interlocks.

A Theory of Strategic Information Technology Unavailability (p. 1571)

Amin K. Amiri, Hasan Cavusoglu, Izak Benbasat,
Varun Grover

This study develops a practice-relevant theory explaining why some information technology (IT) unavailability incidents lead to severe and prolonged organizational consequences. By redefining IT unavailability as unmet demand for IT resources rather than simple system downtime, the paper shows how disruptions cascade through information capacity deficits and business capacity deficits, ultimately impairing critical services, affecting clients, and damaging organizational reputation. An analysis of 28 real-world IT unavailability incidents reveals that three reinforcing feedback loops (IT inertia, information inertia, and business inertia) can intensify service disruption, delay recovery, and amplify downstream impacts.

NFT Disruption in Platform Competition: Evidence from Trading Card Collectibles (p. 1596)

Ioannis Filippou Kanellopoulos, Dominik Gutt, Ting Li

The rise of blockchain-based platforms is reshaping how digital and physical products interact, with important consequences for platform strategy and market regulation. Evidence from the introduction of NBA Top Shot shows that digital collectibles can directly reduce prices, sales volume, and market value in adjacent physical markets, indicating a clear cannibalization effect. However, the impact is highly uneven: Physical products with close digital substitutes and low digital scarcity become especially vulnerable, whereas low-cost items in certain segments can experience market expansion as new collectors enter through digital channels. These patterns demonstrate that digital transformation alters competition not only between platforms but also across product categories, making strategic decisions around scarcity, product design, and release timing critical for firms managing multiformat ecosystems. For practice, the findings highlight the need for companies to anticipate cross-market spillovers when launching digital offerings and to manage physical and digital product lines in an integrated way. For policy, the results underscore the importance of establishing transparent standards for digital collectibles and providing clear consumer guidance, while recognizing that many nonfungible tokens function as collectible goods rather than financial securities and may warrant distinct regulatory treatment.

The Economics of Password Sharing (p. 1621)

Jianqing Chen, Jianchao Sheng

Password sharing is widespread across subscription-based industries, such as streaming services like Netflix. Although sharing reduces out-of-pocket costs for users, it also creates sharing-related burdens, such as privacy concerns and coordination hassles. We develop a game-theoretic model in which a firm serves two consumer segments—password sharers and nonsharers—to analyze how password sharing affects firm performance and consumer outcomes. Our results show that sharing enables groups of users to “self-bundle” by aggregating their valuations and subscribing through a single shared account. It also creates implicit price discrimination; sharers self-select between sharing and subscribing individually under a uniform price. These mechanisms can increase a firm’s revenue when sharing costs are low. In addition, password sharing reshapes the disparity between sharers and nonsharers. When nonsharers have sufficiently higher willingness to pay, the firm is even more likely to benefit from sharing. However, password sharing poses policy concerns. We find that it reduces social welfare and consumer surplus across a broad range of conditions—especially when sharing costs are moderate. Moreover, the firm’s profit-maximizing

strategy, particularly when it chooses to accommodate sharing, often conflicts with what is best for consumers or society. These insights highlight the need for firms to carefully evaluate sharing policies and in some cases, for regulators to intervene to protect consumer welfare.

Workflow Automation in Open-Source Software Development: Accelerating Innovation Through Mechanization and Orchestration (p. 1643)

Ao Huang, Ni Huang, Yili Hong

This study develops a conceptual framework distinguishing two mechanisms of workflow automation: mechanization and orchestration. Mechanization automates discrete, self-contained, repeatable tasks through standardized execution to enhance consistency, reliability, and efficiency, while orchestration automates the communication between tasks, workers, and stages, which facilitates information flow and coordination. We theorize that these mechanisms differentially affect incremental versus substantive innovation. Using a multimethod approach integrating machine learning and econometrics, we analyze the effects of workflow automation in open-source software development, demonstrating that mechanization accelerates maintenance-oriented exploitative innovation, whereas orchestration accelerates development-oriented explorative innovation. This mechanization-orchestration distinction extends beyond software contexts. For practitioners, aligning automation strategies with innovation goals is essential: deploy mechanization to enhance operational efficiency and support incremental improvements in stable environments; and implement orchestration to enable adaptive coordination in exploratory, high-velocity development requiring creativity and flexibility. For policymakers, understanding this distinction informs workforce development and technology adoption policies, as automation reshapes work by shifting human contribution from routine execution toward coordinated problem solving and strategic decision making.

From Shield to Sword: How Data Privacy Can Undermine Data Security (p. 1663)

Alexander Gladis, Torsten-Oliver Salge, David Antons, Nicole Hartwich

What is the point in hacking computer systems when organizations voluntarily disclose personal data to anyone who asks convincingly? We show that the European General Data Protection Regulation (GDPR) is paradoxically exploitable for identity theft despite being designed to protect personal data. Subject access requests (SARs) according to its “right of access” (Article 15) can be weaponized by impersonating a victim and submitting fraudulent SARs in their name. We task attackers with stealing the personal data of three volunteers (highly privacy aware person, average user, and semipublic figure) in a real-world setting. These

attacks could be replicated by just about anyone. Yet, they obtained sensitive personal data, including addresses, phone numbers, national ID and bank account information, and insurance data. Based on 718 submitted SARs and 21 interviews with data protection officers, we tell a frightening, yet fascinating story of how these identity thefts unfold, expose flaws in how organizations process SARs, and uncover a systemic weakness in the GDPR. We analyze the underlying factors enabling such attacks, assess their real-world impact, and explore mitigation options for individuals, organizations, and lawmakers. Our insights have important implications for how data privacy and data security interrelate and how we manage and regulate them.

Learning from Earnings Calls: Graph-Based Conversational Modeling for Financial Prediction (p. 1681)

Yi Yang, Yixuan Tang, Yangyang Fan, Kunpeng Zhang

Earnings conference calls are a critical channel through which public companies communicate with investors, analysts, and regulators. These conversations contain timely signals about firms' future risk, yet their length and unstructured nature make systematic analysis difficult in practice. This study develops an artificial intelligence (AI)-based approach, motivated by a body of theoretical and empirical work from finance and accounting, that transforms earnings call transcripts into structured representations, allowing key aspects of managerial communication, such as topic flow, cross-referencing, sentiment, and other semantics, to be explicitly modeled. Empirical results show that the proposed model significantly improves the prediction of future financial risk compared with existing deep learning and large language model approaches, particularly for firms with complex and lengthy disclosures. The findings highlight that nuanced manager-analyst interactions within earnings calls contain value-relevant information for market participants. In particular, cross-referencing, the introduction of new topics, and a more positive tone are associated with lower subsequent risk as identified by the proposed approach. For researchers and practitioners, this work demonstrates how theoretical and empirical evidence on managerial communication can be incorporated into predictive model design, supporting more accurate, interpretable, and responsible use of AI in financial markets.

Analyzing Consumer Footprints on E-Commerce Platforms: A Multichannel Sequential Search Model with Reference Price (p. 1706)

Hao Zhang, Zhiling Guo, Junming Liu, Mingzheng Wang

Accompanied by the popularity of mini-programs, consumers can easily browse products across both the e-commerce native app channel and the mini-program channel, leaving behind extensive click-through and

purchase records across various channels. These multichannel footprints contain rich information about consumers' search preferences, which offer great potential for optimizing platforms' multichannel management. However, prior studies predominantly focus on single-channel contexts, which may not directly apply to analyzing consumers' multichannel footprints. This research develops a multichannel sequential search (MSS) model to characterize consumers' multichannel footprints as a threefold search process consisting of cross-channel, cross-product, and cross-page search, where consumers dynamically update their cross-channel reference prices as they navigate their search journeys. The estimation results of the MSS model reveal that consumers exhibit significant heterogeneity in channel preferences, cross-channel costs, and sensitivity to cross-channel reference prices, contributing to their diverse multichannel footprints. Drawing on a comprehensive understanding of the MSS process, our optimal policy recommends that the platform implement a promotion gap with an average discount percentage of 32% between the app and mini-program channels. The channel-specific promotion strategy fosters a strong cross-channel reference price effect, which uses a smaller promotion to achieve a significant profit improvement over the state-of-the-art model.

Content Exclusivity on Advertising Revenue-Sharing Platforms (p. 1725)

Yuansheng Wei, Hang Wei, Lin Tian, Baojun Jiang

Digital content platforms such as YouTube, TikTok, and Twitch rely on third-party providers and share advertising revenue to incentivize content production. Yet, as providers increasingly distribute content across multiple platforms, competition intensifies and platforms consider exclusive contracts to restrict providers' multi-homing behavior. Our research develops an analytical model to examine when exclusive contracts benefit or harm platforms, providers, consumers, and overall welfare. We find that exclusivity has two key forces: it limits the number of providers on each platform, reducing same-side crowding, but it also weakens the effectiveness of revenue sharing as a competitive tool for attracting providers. As a result, exclusivity benefits platforms only when consumers place low marginal value on content, leading platforms to offer lower sharing rates and earn higher profits. When consumer valuation is high, exclusivity can backfire, reducing platform profits. Importantly, exclusivity can increase provider surplus and even generate win-win outcomes for both platforms and providers at intermediate valuation levels. Although exclusivity reduces consumer surplus because of fewer providers, total social welfare may still increase because of reduced competition costs among providers. These findings suggest policymakers should take a nuanced stance toward regulating content exclusivity.

A User Purchase Motivation-Aware Product Recommender System (p. 1740)

Jiarong Xu, Jiaan Wang, Hongzhe Zhang, Tian Lu

Retailers struggle to match recommendations to why customers buy. We introduce a practical framework that distinguishes two core, actionable purchase motivations, *stable preference* and *exploratory intent*, and present STB, a data-efficient measure that infers which motivation drives each item purchase using only transaction sequences and item attributes. Building on STB, we develop UPSTAR, a motivation-aware recommender that separates users' behavior into stable-preference and exploratory subsequences and fuses their signals for next-item prediction. Across three real-world e-commerce data sets, UPSTAR substantially improves accuracy and, importantly, advances the system's ability to surface genuinely exploratory items that drive discovery and cross-category sales. For practitioners, our method enables more targeted marketing: promote reliable items to preference-driven buyers while exposing exploratory buyers to curated novelty, improving conversion and long-term engagement. For platform policy and operations, motivation-aware recommendations support inventory planning, personalized promotions, and responsible diversification of exposure without requiring surveys or extensive auxiliary data. Implementation requires only existing transaction logs and item metadata, making it immediately deployable for large-scale retail systems.

Human-Algorithm Collaboration in Gig Work: The Role of Experience, Skill Level, and Task Complexity (p. 1760)

Benjamin Knight, Dmitry Mitrofanov, Serguei Netessine

In this paper, we contribute to recent studies on human-algorithm collaboration by examining how experience, skill level, workload, and task complexity shape the impact of an algorithm-enabled decision-support tool for gig workers. We leverage a large-scale randomized field experiment on the Instacart platform from June 2022 to September 2022. The algorithm-enabled technology aims to revolutionize item picking by helping shoppers locate and collect items more efficiently, reducing picking time while maintaining service quality, as reflected by refund rates. We find that the technology complements experience: rather than diminishing the value of experience, it yields larger improvements for more experienced shoppers. We also find that it substitutes for skill levels by helping lower-skilled workers bridge the performance gap with higher-skilled peers, but lower-skilled workers need experience to fully benefit from the tool. Finally, treatment effects vary with workload and task complexity, clarifying when algorithmic guidance is most valuable. For

policymakers, our findings suggest a simple rule: give workers some baseline experience before introducing artificial intelligence tools, using a staggered rollout with basic training. We also show that these tools can make service more consistent by closing the gap between high performers and lower performers, reducing performance dispersion, and helping standardize quality.

Disclosure of Cybersecurity Investments and the Cost of Capital (p. 1782)

Taha Havakhor, Mohammad S. Rahman, Tianjian Zhang

Executives often hesitate to disclose their company's cybersecurity investments, fearing lawsuits or negative reactions from investors. Our research shows that transparency in this area actually pays off. Analyzing U.S. Securities and Exchange Commission filings from nearly 2,000 public firms, we find that companies that disclose cybersecurity investments enjoy a lower cost of capital—cheaper access to debt and equity financing. These benefits are strongest when disclosures are specific rather than boilerplate, when the firm is followed by more analysts and more institutional investors, and when disclosed investments in cybersecurity are substantial. The takeaway is clear: Meaningful disclosure builds trust with investors, who reward transparency by lowering financing costs. For leaders and regulators alike, this finding highlights that openness about cybersecurity readiness is not just good governance—It is smart business in a capital market that increasingly values risk management and resilience.

SUVA: A Probabilistic Framework for Auditing LLMs with an Application to Social Preferences (p. 1805)

Yan Leng, Yuan Yuan

Organizations are increasingly deploying large language models (LLMs) as customer service agents, decision aids, and semiautonomous agents. We develop State-Understanding-Value-Action (SUVA), a probabilistic auditing framework that turns an LLM's response into structured evidence about how its decision was produced. SUVA treats the prompt as the state, codes the model's reasoning to extract its understanding and stated values using a transparent codebook, and then estimates how these elements statistically predict the eventual action. We demonstrate SUVA on social preference games from behavioral economics and show how the same workflow can audit other delegated decisions by using domain-specific prompts and value codebooks. Across eight widely used LLMs, SUVA reveals systematic prosocial and reciprocity patterns and shows how posttraining alignment reshapes them. For practice and policy, SUVA supports a repeatable audit, align, and reaudit workflow for model selection, compliance, and ongoing monitoring of deployed LLM systems.

Conform or Workaround? A Multilevel Analysis of the Effect of Group Cultural Tightness on Enterprise System Use (p. 1831)

Shaobo Wei, Xiayu Chen, Ronald E. Rice, Chee-Wee Tan, Yezheng Liu

Enterprise systems (ESs) embed industrial best practices into adopting organizations through their system features, but employees frequently “work around” the system to get work done—sometimes in helpful ways, sometimes in risky ones. Our research shows that group cultural tightness (strongly enforced norms within a team) is a powerful lever for ES governance: tighter groups reliably increase conforming use while reducing both internal and external workarounds, and this pattern holds across Chinese and U.S. contexts and multiple organizational settings. Importantly, not all workarounds are equal. Evidence from multisourced, longitudinal field data indicates that conforming use and internal workarounds can improve job performance, whereas external workarounds harm performance—likely by creating fragmented processes, data gaps, and compliance and security exposure. For practice and policy, the message is clear: organizations should strengthen team-level norms and accountability to curb harmful “outside-the-system” behaviors, while simultaneously creating safe, sanctioned channels to surface, evaluate, and integrate beneficial internal workarounds (e.g., controlled extensions, approved templates, and rapid governance reviews). This balanced approach supports performance, compliance, and continuous improvement without shutting down frontline problem-solving.

Are You, You? Seamlessly Fighting Identity Fraud with Keystroke Dynamics (p. 1854)

David Kim, Joseph S. Valacich, Jeffrey L. Jenkins, David W. Wilson, Manasvi Kumar, Paul Weisgarber

In September of 2017, Equifax disclosed a data breach that exposed the personal information of 147 million people, including names, Social Security numbers, and addresses. Such high-profile data breaches have rendered traditional forms of identity verification—especially knowledge-based authentication (KBA)—worse than useless: fraudsters have a 92% success rate in KBA screenings, compared to just 46% for genuine customers. In the face of these challenges, digital platforms are turning to sparse alternative data sources and overt verification technologies, often to the detriment of the user experience. The objective of this research is to design and build a novel approach to identity verification for new platform users using digital behavior data—features that describe how users type and interact during account setup. The system (1) evaluates identity fraud risk for all first-time users, and (2) minimizes the impact on the new user experience by

seamlessly analyzing behavior during a platform’s existing onboarding experience. We evaluated and improved the design in four experiments, culminating in an identity fraud detection tool that effectively detects identity fraud for first-time users and supports seamless user experiences.

AI Governance and the Decentralization of Technology Production: An Investigation of AI-Based IPA Bots (p. 1873)

Abhishek Kathuria, Prasanna P. Karhade, Ojaswi Malik, V. K. Pani Baruri, Benn R. Konsynski

We revisit the centralization–decentralization tension in the context of decentralized technology production at the artificial intelligence (AI) frontier, focusing on Intelligent Process Automation (IPA) bots as a salient manifestation of the democratization of AI. IPA bots combine robotic process automation with AI technologies and process mining based on deep, mindful domain expertise. We collaborate with a Fortune 200 multinational to study how IPA projects yield successful governance outcomes of utilization and repeatability. Our research reveals that traditional centralized mandates, when paired with the unique learning and adaptive capabilities of AI systems, can actively suppress utilization and stifle organizational reuse of the bots. Conversely, democratization is no panacea. While empowerment of business users through decentralization usually leads to successful outcomes for the deployed bots, when certain boundary conditions are crossed, these same decentralized efforts can produce fragmentation in the form of one-off solutions. Our research challenges core tenets of the information technology governance canon, demonstrating that centralization–decentralization logic at the AI frontier must expand to account for the evolution of governance, process, and technological characteristics to realize effective AI governance.

Reference Aware Delexicalization (RAD) Framework: Theory Driven Artificial Intelligence Modeling for Domain Generalization (p. 1895)

Sandeep Suntwal, Susan A. Brown

Artificial intelligence systems often underperform in new contexts, limiting their reliability in business domains. This study introduces the reference aware delexicalization (RAD) framework, a theory-driven approach that improves artificial intelligence (AI) model performance across different domains without requiring massive computational resources. RAD addresses a fundamental problem: AI models often memorize surface patterns from training data rather than learning transferable reasoning skills. By systematically abstracting domain-specific terms while preserving semantic relationships, RAD enables models to focus on underlying logical structures that generalize across contexts. For practitioners, RAD offers measurable benefits. Healthcare

organizations can deploy clinical decision support systems that maintain accuracy when processing records from different departments or institutions. Financial institutions can build fraud detection systems that adapt to emerging threats without extensive retraining. Social media platforms can improve content moderation consistency across languages and cultural contexts. For policymakers, RAD demonstrates that effective AI does not require ever-larger models with corresponding environmental and economic costs. Organizations can achieve robust, adaptable AI systems through principled data augmentation techniques. This finding supports policies encouraging efficient, interpretable AI development over resource-intensive scaling approaches, promoting both technological sustainability and broader access to reliable AI capabilities.

The Differential Diffusion of Exchange and Utility Value Blockchain Tokens (p. 1920)

Yegin Genc, Harris Kyriakou, Likoebe Mohau Maruping, Ling Xue

Blockchain tokens are widely used in digital markets, yet they do not spread in the same way. Our research shows that tokens designed primarily for financial exchange and those designed for functional or service use diffuse through fundamentally different mechanisms even though they operate on the same underlying technology. Analyzing more than 200 million transactions across more than 24,000 Ethereum-based tokens, we find that financial tokens spread faster when they are held as part of large, diversified user portfolios yet diffuse more slowly when adopted by highly visible lead users. In direct contrast, utility tokens benefit strongly from early adoption by lead users and from transaction platforms that facilitate repeated use, whereas portfolio diversification inhibits their diffusion. These findings have important implications for practice and policy. For managers and platform designers, strategies that successfully promote financial tokens may actively hinder the adoption of utility tokens, underscoring the need for value-specific design, marketing, and platform strategies. For regulators and policymakers, the results provide empirical evidence that blockchain tokens function differently depending on their value orientation, constituting a strong empirical basis for more nuanced, behavior-based regulatory frameworks rather than one-size-fits-all approaches to digital tokens.

To Collect or to Purchase? Collaboration Between Manufacturer and E-Commerce Platform on Customization (p. 1944)

Ping Yan, Jun Pei, Subodha Kumar

More manufacturers are combining novel customer-to-manufacturer (C2M) with traditional make-to-order

(MTO) customization to deliver tailored products. However, they face a critical choice between collecting data independently and purchasing data-driven insights from e-commerce platforms. This key decision remains ambiguous in practice and understudied in the information systems literature, with little rigorous analysis on data acquisition ways and their interaction with hybrid customization. This study addresses this gap by investigating firms' data acquisition strategies in the context of hybrid customization. We find that when the marginal cost of data acquisition is high for both the manufacturer and the e-commerce platform, and customers have a stronger bias against tailored products, the manufacturer should still prefer collecting market data independently rather than purchasing data-driven insights. The e-commerce platform should not always provide insight support. Moreover, manufacturer-collected C2M may sometimes lead to a higher effort level of acquiring informative data and more customer welfare than platform-initiated C2M. These results provide valuable guidelines to manufacturers on choosing an appropriate data acquisition way, to e-commerce platforms on offering data-driven insights, and to policymakers on improving customer welfare. These results also contribute to the emerging literature in IS and related disciplines on data-driven production models.

Contribute to MY IT Service: Encouraging Technology Extra-Role Behaviors in User-Artifact Interactions from a Psychological Ownership Perspective (p. 1966)

Haiyun (Melody) Zou, Yulin Fang, Heshan Sun, Kai H. Lim

This research aims to investigate how interactions with information technology (IT) artifacts can sustain users' continuance usage of, and encourage their voluntary contributions to, access-based and algorithm-empowered IT services (e.g., Spotify and Duolingo). This question is particularly important because in such contexts, users often have limited interactions with other human users—traditionally the main driver of the target behaviors in virtual communities. In the results, we found that users are more willing to use the technology and perform extra-role behaviors that benefit the technology when they develop a user-artifact relationship, after controlling the alternative mechanisms. We further explain how customization and personalization features facilitate users' development of such a relationship by co-constructing an extended self with the IT artifact. For practitioners, our study mobilizes the user resources and promotes user voluntary contributions that benefit the technology (e.g., knowledge contribution, user feedback, new user referral, and voluntary payment). By incorporating the algorithm into self-concept as an extended self, we propose a new way of human-algorithm interaction, and hence are able to advise on algorithm appreciation and AI adoption in the

non-competitive consumer technology context and provide guidelines and use cases on IT design regarding customization and personalization.

Beyond Truthful Reporting: Robust Strategies for Worst-Case Payoff Maximization in Large Markets (p. 1994)

Saša Pekeč, Chenxi Xu

In many large-scale markets mediated by digital platforms, such as advertising, energy trading, transportation logistics, and spectrum auctions, platforms suggest that participants report true valuations to simplify strategic complexity. This recommendation is widely followed by smaller bidders who lack resources for sophisticated strategic analysis. This paper shows that simple deviations from truthful reporting can improve outcomes for such participants. Using a robust optimization framework, we derive practical bidding strategies for generalized first-price, generalized second-price, and core-selecting combinatorial auctions. Our distribution-free approach maximizes worst-case payoffs over a set of plausible competitor bids, without requiring detailed information about rival behavior. We demonstrate that simple strategies, specifically shading bids on preferred bundles, consistently outperform truthful bidding when bidders have straightforward valuation structures. This insight persists even in complex settings where closed-form optimal policies cannot be derived. The results provide actionable guidance for market participants with limited information who lack resources for deploying sophisticated bidding algorithms, while offering market designers insight into how participants might deviate from recommended strategies.

Balancing Acts: Unveiling the Dynamics of Post Removal on Social Media User-Generated Content (p. 2012)

Guohou Shan, Liangfei Qiu

User-generated content (UGC) is central to engagement and value creation on social media platforms, but policy violations and low-quality posts create significant reputational, legal, and financial risks. Platforms commonly respond by removing posts that violate community standards, yet it remains unclear whether such actions deter users from contributing or instead improve subsequent behavior. Analyzing data from 40 Reddit communities using a difference-in-differences design, we examine how post removal influences users' future contributions. We find that removing a user's post reduces subsequent rule violations and increases the average number of upvotes their later posts receive, an indicator of improved content quality. These effects are stronger when users experience repeated removals and when they observe peers' posts being removed, highlighting both direct and vicarious learning mechanisms. For platform operators, the findings suggest that targeted,

consistent, and transparent post removal can serve as an effective governance tool, not merely to suppress harmful content but to foster higher-quality participation over time. For policymakers and regulators, the results underscore the importance of moderation frameworks that promote user learning, accountability, and procedural fairness. Well-designed moderation systems can enhance community standards while maintaining legitimacy and trust in digital governance.

Bargaining over Data and Analytics: Sellers, Buyers and Consultants (p. 2031)

Jyotishka Ray, Syam Menon, Vijay Mookerjee

Most firms routinely gather vast amounts of data as part of doing business. Monetizing this proprietary data is an increasingly attractive revenue stream for many of these firms. Two fundamental decisions need to be made when deciding to sell, the first involving a choice between selling exclusively to a single buyer or nonexclusively to many, and the second related to what exactly should be sold—just the data or a data product that bundles data with analytics services. As firms often resort to bargaining to arrive at sale agreements, we analyze these decisions through the lens of a bargaining framework. When the buyer needs help from a third-party consultant, she too needs to make a decision—between negotiating separately with the seller and the consultant, and with both of them simultaneously. We find that there are situations where sellers can benefit from bundling the data with analytics services even when their analytics capabilities are weak relative to those of external consultants. Simultaneous negotiations enable buyers to extract more of the consultant's contribution, making them preferable to buyers when consultants add substantial value. Broadly, this study provides a road map for structuring contracts to firms considering the sale of their proprietary data.

Forget Me If You Can: Auditing User Data Revocation in Recommendation Systems (p. 2049)

Zhihao Zhu, Yi Yang, Yangyang Fan, Defu Lian

Personalized recommendation systems power today's e-commerce and streaming platforms, but they also raise regulatory concerns about user privacy. Under regulations such as the General Data Protection Regulation (GDPR), individuals have the "right to be forgotten." Whereas companies can remove user records from databases, it remains unclear whether trained artificial intelligence (AI) models truly forget the behavioral patterns or preferences learned from that data. If not, platforms may continue profiling users even after deletion requests. This study introduces RecAudit, a practical auditing tool that tests whether a recommender system has genuinely removed a user's influence from its trained model. RecAudit acts as an independent verification mechanism that identifies users whose behavioral traces persist after

data deletion. Across multiple real-world data sets, RecAudit substantially outperforms existing auditing and membership inference methods. By enabling organizations to detect high-risk cases and target corrective machine unlearning efforts, RecAudit provides a concrete technical pathway to operationalize data revocation rights. The framework supports platform operators and regulators in strengthening accountability and ensuring AI systems comply with evolving data protection laws.

Leveraging Multiview Data Through Discrete and Regularized Deep Learning for Dynamic Financial Risk Prediction (p. 2068)

Zhao Wang, Wanliu Che, Cuiqing Jiang, Huimin Zhao

Given the dramatic surge of demand for predictive insights into the dynamics of financial risk and the rich, yet entangled, information brought by proliferating multiview data, we propose a discrete and regularized deep learning (DRDL) method to better leverage such multiview data for dynamic financial risk prediction. Empirical evaluation demonstrates advantages of DRDL over benchmarked classic and state-of-the-art methods at both the model level (time-to-risk and out-of-time prediction performance) and the application level (identification and profitability performance). Besides performance gains, DRDL offers distinctive practical advantages. First, it enables explicit and controllable factor-level representations, allowing practitioners to inspect and regulate how cross-view signals are encoded. Second, it offers unique advantages in explicitly and precisely filtering out redundant information while extracting complementary information across heterogeneous data sources, allowing practitioners to better understand which unique informational components drive risk predictions. Third, it offers a practically viable and empirically effective way to promote functional disentanglement within a discrete and structured latent space. Fourth, it supports both time-wise and instance-wise monotonicity, aligning predictions with the cumulative and irreversible nature of financial risk escalation, which may be particularly valuable in risk monitoring and governance contexts.

More Can Be Less: The Economics of Answer Viewing on Paid Q&A Platforms (p. 2091)

Yi Gao, Amit Mehra, Dengpan Liu

Paid Q&A platforms increasingly experiment with an *answer-viewing* feature that allows users to pay a small fee to access answers already given to others. At first glance, this model appears to expand access and create

new revenue streams. Our research, however, shows that the reality is more complex. Using a game-theoretic model, we demonstrate that answer viewing can sometimes *harm* platform profits due to market cannibalization: Users who might have paid for personalized answers switch to cheaper viewing options. Surprisingly, we find that answerers may raise their consulting fees when platforms share a large portion of viewing revenue with them, and users do not always benefit because personalized answers become more expensive. These results reveal a central paradox—*more can be less*—as adding another monetization channel may backfire. For practice, our findings highlight the need for Q&A platforms to carefully calibrate revenue-sharing policies to avoid undermining their core market. For policy, the study underscores the importance of safeguards to protect small knowledge providers when platform incentives diverge from those of answerers and users. Beyond Q&A, the insights apply to broader contexts where platforms repackage personalized services into standardized products.

Where the Ball Starts Rolling? An Empirical Investigation into Initial Opinion Formation on Social Media Platforms (p. 2113)

Venu Puthineedi, Ashish Kumar Jha

Social media users routinely encounter unfamiliar, high-stakes information and form impressions that shape what they believe, share, and act on. This article shows that these first impressions can stabilize after only a small, bounded run of consistent exposure, creating an early “sufficiency” point after which additional posts add little. Once this early stance forms, it guides downstream engagement: users are more receptive to information that aligns with their initial view and less responsive to information that challenges it. Importantly for practice and policy, early judgments are often weakly tied to factual accuracy under low-effort scrolling, making it difficult for later corrections to fully reverse the effects of early exposure. The study also demonstrates that source cues matter: professional titles and affiliations can outweigh visible platform signals such as badges and popularity metrics when people evaluate unfamiliar content. Together, these findings suggest that platforms and regulators should prioritize upstream interventions—verifying expertise claims before allowing professional titles to be displayed, strengthening domain-specific credibility signals, and ensuring authoritative information reaches users early in exposure sequences—especially in sensitive areas such as health, finance, and public safety.