



International Journal of Research Publication and Reviews

Journal Homepage: www.ijrpr.com ISSN 2582-7421

The Influence of AI-Enabled Smart Recommendations on Consumer Satisfaction and Choice Architecture

Mr. Sachin Vasant Chaugule¹, Dr. Bhalchandra Bite²

¹Research Scholar, ²Research Guide

Research Centre: AISSMS Institute of Management (MBA), Pune – 411001, India

ABSTRACT

Purpose: This study examines how AI-enabled smart recommendation systems (AI-SR) affect consumer satisfaction and the way digital platform's structure consumer choices. It focuses on four perceived AI-SR characteristics—personalisation accuracy, algorithmic transparency, contextual relevance, and recommendation diversity—and examines consumer trust and perceived autonomy preservation as explanatory mechanisms. Privacy risk perception and digital literacy are considered as boundary conditions. The study proposes and tests the AI Recommendation Satisfaction and Choice Architecture Model (ARSCAM).

Design/methodology/approach: A quantitative, cross-sectional survey was conducted among 412 adult users of AI recommendation systems in India. Respondents were drawn from e-commerce, streaming/entertainment, and financial-technology contexts. Data were analysed using SPSS 27 and SmartPLS 4.1.

Findings: Personalisation accuracy showed significant positive relationships with utilitarian, hedonic, and relational satisfaction. Transparency was positively related to trust and satisfaction. Contextual relevance predicted satisfaction, while recommendation diversity predicted hedonic satisfaction and consideration-set breadth. AI-SR significantly affected option framing, default structuring, and consideration-set composition. Trust and perceived autonomy preservation played mediating roles. Privacy risk weakened the personalisation–satisfaction relationship, whereas digital literacy strengthened the transparency–trust relationship. The model explained 71.2% of the variance in consumer satisfaction and 64.8% in choice architecture.

Practical implications: Recommendation platforms can improve consumer experience by combining accurate personalisation with understandable explanations, privacy controls, diverse recommendations, and mechanisms that allow users to expand or modify the algorithmically curated option space.

Keywords: AI recommendation systems; consumer satisfaction; choice architecture; personalisation; transparency; trust; autonomy; privacy risk; digital literacy; PLS-SEM.

1. Introduction

Digital platforms increasingly mediate consumer decisions through algorithms that filter, rank, and present information. The growth of product catalogues, streaming libraries, and digital financial services makes unrestricted search costly in time and attention. Recommendation systems address this problem by narrowing the information environment and identifying options that are predicted to be relevant. Their role, however, extends beyond convenience: by determining what is prominently displayed, what is presented as a default, and which alternatives enter the consumer's active consideration set, recommendation systems also participate in choice architecture.

The literature describes recommendation quality through characteristics such as accuracy, relevance, diversity, explainability, and personalisation. Accurate recommendations can reduce search effort and improve task performance, while diverse recommendations can encourage exploration and serendipitous discovery. At the same time, highly personalised systems may create privacy concerns when users infer that extensive behavioural data are being collected. Limited explanation of algorithmic logic may further reduce confidence in the recommendations. These tensions make satisfaction a multidimensional outcome rather than a simple consequence of recommendation accuracy.

Choice architecture is particularly relevant because AI systems dynamically organise decision environments. Traditional choice architectures are usually designed in advance, whereas AI-mediated environments can adapt to an individual's history, context, and responses. Such systems may support consumers by simplifying complex decisions, but they can also narrow perceived alternatives or shape attention through defaults and framing. The present research therefore treats consumer satisfaction and choice architecture as related but distinct outcomes.

India provides an important setting for examining these relationships because digital consumption spans heterogeneous consumer groups and multiple platform categories. The original study focuses on five metropolitan markets—Mumbai, Delhi-NCR, Bengaluru, Hyderabad, and Ahmedabad—and three broad digital contexts: e-commerce, streaming/entertainment, and financial technology.

The study addresses three principal gaps identified in the original paper. First, the four major AI-SR characteristics are examined within one structural framework rather than in isolation. Second, choice architecture is operationalised through option framing, default structuring, and consideration-set composition. Third, the study provides empirical evidence from Indian consumers and integrates trust, perceived autonomy, privacy risk, and digital literacy into the same model.

2. Literature Review and Theoretical Foundation

2.1 Technology Acceptance Model

The Technology Acceptance Model (TAM) explains technology-related behaviour through perceived usefulness and perceived ease of use (Davis, 1989). In recommendation environments, usefulness can be reflected in the extent to which recommendations match consumer needs and context, while ease of use can be reflected in interface simplicity and understandable explanations. The present study uses this logic to connect personalisation accuracy, contextual relevance, and transparency with consumer evaluations.

2.2 Self-Determination Theory

Self-Determination Theory (Deci & Ryan, 1985) emphasises autonomy, competence, and relatedness. In digital recommendation settings, autonomy concerns whether users feel that the system assists rather than replaces their decisions. A recommendation may be accepted when it reduces effort while preserving the consumer's ability to explore, reject, reorder, or customise the available alternatives. Perceived autonomy preservation is therefore incorporated as a mediator between AI-SR and choice-architecture outcomes.

2.3 Nudge Theory and Choice Architecture

Choice architecture concerns how alternatives are organised and presented without necessarily removing options or changing economic incentives (Thaler & Sunstein, 2008; Johnson et al., 2012). Digital recommendation systems can influence attention through ranking, framing, defaults, social signals, and the size and composition of the consideration set. This perspective provides the conceptual basis for examining AI-SR as an adaptive form of choice architecture.

2.4 Expectation-Confirmation and Satisfaction

Expectation-Confirmation Theory links satisfaction to the relationship between expected and experienced performance (Oliver, 1980; Bhattacharjee, 2001). For AI recommendations, satisfaction is likely to increase when recommendations meet expectations for relevance, accuracy, usefulness, and ease of decision making. The present study distinguishes utilitarian, hedonic, and relational dimensions of satisfaction so that functional value, enjoyment/discovery, and perceived personal understanding are not treated as identical outcomes.

2.5 AI-SR Characteristics

Personalisation accuracy reflects how closely recommendations correspond to a user's preferences, needs, and situation. Prior research associates perceived personalisation with trust and adoption, while privacy research identifies a personalisation-privacy tension in which unusually accurate recommendations may appear intrusive (Awad & Krishnan, 2006; Komiak & Benbasat, 2006).

Algorithmic transparency concerns the extent to which users can understand why an item was recommended. Explanation mechanisms such as purchase-history or similarity-based rationales can make recommendations more interpretable and can strengthen confidence in system competence (Sinha & Swearingen, 2002; Tintarev & Masthoff, 2007; Gedikli et al., 2014).

Contextual relevance concerns whether the recommendation fits the immediate task, situation, or consumption context. Recommendation diversity concerns whether the system presents sufficiently varied alternatives rather than repeatedly narrowing users toward established preferences. Diversity and serendipity complement accuracy by supporting exploration and discovery (Zhang et al., 2012; Kotkov et al., 2016; McNee et al., 2006).

2.6 Research Gap

The reviewed literature indicates strong evidence for individual recommendation characteristics, but fewer studies combine these characteristics with satisfaction and choice architecture in one empirical model. The present study therefore integrates the four AI-SR dimensions with trust, autonomy, privacy risk, and digital literacy and tests their relationships simultaneously.

3. Research Problem, Objectives and Hypotheses

The central research problem is: How do characteristics of AI-enabled smart recommendation systems influence consumer satisfaction and choice architecture, through which mediating mechanisms, and under what moderating conditions among Indian digital consumers?

3.1 Objectives

- To validate the principal consumer-perceived characteristics of AI-enabled smart recommendation systems.
- To examine the effects of personalisation accuracy, transparency, contextual relevance, and diversity on consumer satisfaction.
- To assess AI-SR effects on option framing, default structuring, and consideration-set composition.
- To test consumer trust and perceived autonomy preservation as mediators.
- To examine privacy risk perception and digital literacy as moderators.
- To validate the integrated AI Recommendation Satisfaction and Choice Architecture Model (ARSCAM).

3.2 Hypotheses

Code	Hypothesis
H1a	Personalisation accuracy positively influences utilitarian satisfaction.
H1b	Personalisation accuracy positively influences hedonic satisfaction.
H1c	Personalisation accuracy positively influences relational satisfaction.
H2a	Algorithmic transparency positively influences consumer trust.
H2b	Algorithmic transparency positively influences consumer satisfaction.
H3	Contextual relevance positively influences consumer satisfaction.
H4a	Recommendation diversity positively influences hedonic satisfaction.
H4b	Recommendation diversity positively influences consideration-set breadth.
H5a	AI-SR significantly influences option-framing effects.
H5b	AI-SR significantly influences default structuring.
H5c	AI-SR significantly influences consideration-set composition.
H6a	Consumer trust mediates the AI-SR–satisfaction relationship.
H6b	Perceived autonomy preservation mediates the AI-SR–choice-architecture relationship.
H7a	Privacy risk negatively moderates the personalisation-accuracy–satisfaction relationship.
H7b	Digital literacy positively moderates the transparency–trust relationship.

4. Methodology

4.1 Research Design

The study follows a positivist, hypothetico-deductive approach and uses a quantitative cross-sectional survey. Respondents were adult Indian consumers who had used AI-powered recommendations during the previous three months and could recall relevant interactions. The survey covered e-commerce/online retail, streaming/entertainment, and financial technology.

4.2 Population and Sampling

The sampling frame covered consumers in Mumbai, Delhi-NCR, Bengaluru, Hyderabad, and Ahmedabad. A target of 400 observations was established for PLS-SEM; 412 usable responses were retained from 468 distributed questionnaires. The original study reports 422 returns, ten exclusions, and an effective response rate of 88.0%.

Characteristic	n	%
Male	239	58.0
Female	164	39.8

Characteristic	n	%
Non-binary / prefer not to disclose	9	2.2
18–24 years	96	23.3
25–34 years	154	37.4
35–44 years	102	24.8
45–54 years	42	10.2
55+ years	18	4.3

4.3 Measurement

The questionnaire contained adapted measures from established scales. All substantive items used a seven-point Likert scale from 1 (strongly disagree) to 7 (strongly agree). Content validation involved seven academic experts and three industry practitioners, and a pilot study of 48 respondents was conducted to assess clarity and response burden.

Construct	Items	Principal source(s)
Personalisation accuracy	7	Knijnenburg et al. (2012); Komiak & Benbasat (2006)
Algorithmic transparency	6	Sinha & Swearingen (2002); Tintarev & Masthoff (2007)
Contextual relevance	5	Nelson (1970); Park et al. (2007)
Recommendation diversity	5	Zhang et al. (2012); Kotkov et al. (2016)
Consumer trust	7	McKnight et al. (2002); Gefen et al. (2003)
Perceived autonomy preservation	5	Deci & Ryan (1985); Peters et al. (2018)
Privacy risk perception	5	Malhotra et al. (2004); Smith et al. (1996)
Digital literacy	4	Eshet-Alkalai (2004); Van Dijk (2006)
Consumer satisfaction	10	Oliver (1980); Bhattacharjee (2001)
Choice architecture	9	Johnson et al. (2012); Thaler & Sunstein (2008)

4.4 Data Analysis

Analysis proceeded through descriptive assessment, factor evaluation, measurement-model testing, structural-path estimation, mediation testing, and moderation analysis. SmartPLS 4.1 and SPSS 27.0 were used. Reliability was assessed with Cronbach's alpha and composite reliability; convergent validity with average variance extracted (AVE); and discriminant validity with the Fornell–Larcker criterion and HTMT. Structural paths were tested using 10,000 bootstrap resamples.

5. Results

5.1 Sample and Platform Context

Primary platform context	n	%
E-commerce (Amazon / Flipkart / Myntra)	168	40.8
Streaming (Netflix / YouTube / Spotify)	147	35.7
Financial technology (Zerodha / Groww / Paytm)	97	23.5
Recommendation interaction frequency	n	%
Multiple times per day	118	28.6
Once per day	94	22.8

Primary platform context	n	%
Several times per week	132	32.1
Once per week	48	11.7
Less than once per week	20	4.8

5.2 Common Method Bias

The study used procedural separation of predictor and outcome items, anonymity assurances, and statistical checks. Harman's single-factor test indicated that the largest factor accounted for 24.7% of variance, below the 50% benchmark used in the original analysis. The unmeasured latent common-method-factor check did not materially alter coefficients or significance levels.

5.3 Measurement Model

Construct	A	CR	AVE	Loading range
Personalisation accuracy	0.891	0.914	0.612	0.703–0.864
Algorithmic transparency	0.864	0.900	0.601	0.688–0.841
Contextual relevance	0.847	0.888	0.614	0.714–0.829
Recommendation diversity	0.838	0.882	0.601	0.696–0.833
Consumer trust	0.902	0.921	0.637	0.724–0.871
Perceived autonomy preservation	0.858	0.896	0.628	0.718–0.847
Privacy risk perception	0.841	0.884	0.604	0.694–0.826
Digital literacy	0.814	0.872	0.630	0.731–0.849
Consumer satisfaction	0.921	0.934	0.591	0.672–0.872
Choice architecture	0.908	0.924	0.578	0.661–0.856

All reported alpha and composite-reliability values exceeded 0.75 and 0.82, respectively, and AVE values exceeded 0.50. The original analysis also reported satisfactory Fornell–Larcker and HTMT results, with HTMT values below 0.82.

5.4 Structural Model

Index	Value	Interpretation
SRMR	0.052	Good fit
NFI (PLS)	0.934	Good fit
Chi-square/df	2.07	Good fit
RMSEA	0.051	Good fit
CFI	0.941	Good fit
R ² Consumer satisfaction	0.712	Substantial
R ² Choice architecture	0.648	Substantial
Q ² Consumer satisfaction	0.441	Large predictive relevance
Q ² Choice architecture	0.387	Large predictive relevance

5.5 Hypothesis Testing

Hyp.	Path	β	t	p
H1a	Personalisation $\rightarrow$ utilitarian satisfaction	0.447	8.914	<.001
H1b	Personalisation $\rightarrow$ hedonic satisfaction	0.382	7.218	<.001
H1c	Personalisation $\rightarrow$ relational satisfaction	0.361	6.744	<.001
H2a	Transparency $\rightarrow$ trust	0.429	8.216	<.001
H2b	Transparency $\rightarrow$ satisfaction	0.318	6.002	<.001
H3	Contextual relevance $\rightarrow$ satisfaction	0.361	6.814	<.001
H4a	Diversity $\rightarrow$ hedonic satisfaction	0.274	5.118	<.001
H4b	Diversity $\rightarrow$ consideration-set breadth	0.312	5.847	<.001
H5a	AI-SR $\rightarrow$ option framing	0.391	7.342	<.001
H5b	AI-SR $\rightarrow$ default structuring	0.344	6.411	<.001
H5c	AI-SR $\rightarrow$ consideration-set composition	0.428	8.107	<.001
H6a	Trust $\rightarrow$ satisfaction (mediation)	0.341	6.388	<.001
H6b	Autonomy $\rightarrow$ choice architecture (mediation)	0.298	5.542	<.001
H7a	Personalisation $\times$ privacy risk	-0.187	3.481	<.001
H7b	Transparency $\times$ digital literacy	0.164	3.046	<.01

All hypotheses reported in the original analysis were supported. The confidence intervals excluded zero for the reported structural effects.

5.6 Mediation and Moderation

Mediated relationship	Indirect effect	95% CI	Type
Personalisation $\rightarrow$ trust $\rightarrow$ satisfaction	0.152	0.094–0.214	Partial
Transparency $\rightarrow$ trust $\rightarrow$ satisfaction	0.146	0.089–0.205	Full
Contextual relevance $\rightarrow$ autonomy $\rightarrow$ choice architecture	0.112	0.058–0.172	Partial
Diversity $\rightarrow$ autonomy $\rightarrow$ choice architecture	0.098	0.047–0.154	Partial
Personalisation $\rightarrow$ trust $\rightarrow$ choice architecture	0.119	0.066–0.177	Partial
Transparency $\rightarrow$ trust $\rightarrow$ choice architecture	0.128	0.073–0.186	Full

For privacy risk, the personalisation-to-satisfaction path was weaker among respondents with higher privacy risk perception ($\beta = 0.312$) than among those with lower privacy risk perception ($\beta = 0.478$), with a significant group difference ($p = 0.008$). For digital literacy, the transparency-to-trust relationship was stronger in the higher-literacy group ($\beta = 0.502$) than in the lower-literacy group ($\beta = 0.336$).

6. Discussion

The results indicate that recommendation quality is multidimensional. Personalisation accuracy had positive relationships with all three satisfaction dimensions, with the largest coefficient observed for utilitarian satisfaction ($\beta = 0.447$). This pattern is consistent with the functional role of recommendation systems in reducing search effort and improving task completion. The positive relationships with hedonic and relational satisfaction suggest that accurate recommendations can also contribute to discovery and the perception that a platform understands the consumer.

Transparency was positively associated with trust ($\beta = 0.429$) and satisfaction ($\beta = 0.318$). The mediation results indicate that trust is an important mechanism connecting transparency with consumer outcomes. This finding supports the argument that explanations are not merely interface features; they can help consumers interpret why a recommendation appeared and form judgments about the system's competence and reliability.

Contextual relevance also predicted satisfaction ($\beta = 0.361$). A recommendation can be accurate in general but unsuitable for the user's immediate purpose; contextual fit therefore adds explanatory value beyond generic personalisation. Recommendation diversity showed significant effects on hedonic

satisfaction ($\beta = 0.274$) and consideration-set breadth ($\beta = 0.312$), supporting the view that consumers benefit not only from relevance but also from opportunities for exploration.

The choice-architecture findings show that AI recommendations influence how alternatives are organised. The strongest reported choice-architecture coefficient was for consideration-set composition ($\beta = 0.428$), followed by option framing ($\beta = 0.391$) and default structuring ($\beta = 0.344$). These results are consistent with the basic operation of recommender systems: they influence which alternatives receive attention and therefore which options consumers actively compare.

Perceived autonomy preservation partially mediated the relationship between AI-SR and choice architecture. This suggests that consumers' interpretation of recommendation interfaces matters alongside the interface structure itself. When recommendations are experienced as assistance that can be questioned or modified, their influence may be compatible with a stronger sense of agency.

The moderation results qualify the benefits of personalisation and transparency. Higher privacy risk perception weakened the positive personalisation-satisfaction relationship, indicating that relevance may lose value when users associate it with intrusive data practices. Conversely, greater digital literacy strengthened the relationship between transparency and trust, implying that explanations are more useful when consumers have sufficient ability to interpret them.

7. Managerial Implications

First, platform designers should combine personalisation with privacy communication. Clear information about what data are used, meaningful controls over personalisation, and data-minimisation practices can reduce uncertainty while retaining recommendation utility.

Second, recommendation explanations should be understandable and proportionate to user needs. Short explanations can communicate the principal reason for a recommendation, while optional detail can support users who want to understand the underlying signals more deeply. Explanation interfaces should avoid presenting opaque algorithmic outputs as unquestionable decisions.

Third, choice architecture should preserve practical user control. Features such as preference editing, filtering, reordering, alternative exploration, and clear ways to leave the recommended set can allow consumers to benefit from curation without making the algorithm the sole gatekeeper of available options.

Fourth, recommendation diversity should be evaluated alongside accuracy. Platform evaluation can include measures of discovery, breadth, repeat exposure, and longer-term satisfaction rather than relying exclusively on immediate clicks or conversions.

Fifth, platform context matters. E-commerce systems may emphasise accuracy and contextual fit; streaming services may place greater emphasis on diversity and discovery; and financial platforms may require especially clear explanations and autonomy-supportive choice presentation because users may make higher-consequence decisions.

8. Theoretical Contribution

The ARSCAM framework integrates TAM, Self-Determination Theory, Expectation-Confirmation Theory, and choice-architecture theory in a single empirical structure. Its distinctive feature is the simultaneous treatment of satisfaction and choice architecture as outcomes of AI-SR characteristics, with trust and perceived autonomy as explanatory mechanisms and privacy risk and digital literacy as boundary conditions.

The results provide empirical support for the proposition that recommendation systems are not only prediction tools but also components of consumers' decision environments.

9. Limitations and Future Research

The cross-sectional design limits conclusions about temporal change. Future longitudinal studies could examine whether satisfaction initially rises as recommendation accuracy improves and later changes as recommendation diversity narrows or consumer expectations evolve. The sample is also limited to five Indian metropolitan markets and therefore should not be treated as nationally representative without further evidence.

The study relies on self-reported perceptions, which may not fully capture automatic attention and behavioural responses. Future work could combine survey data with clickstream records, controlled experiments, eye-tracking, or other behavioural measures. Generative-AI recommendation agents also warrant separate investigation because conversational interaction may alter perceived autonomy, explanation quality, and relational satisfaction.

Further research should examine potentially manipulative recommendation designs and the relationship between commercial optimisation, defaults, framing, and consumer welfare. Cross-cultural research is also needed to test whether perceptions of autonomy and delegated choice differ across social and cultural contexts. Finally, legal and consumer-protection research can examine appropriate standards for transparency, explainability, privacy, and autonomy in AI-mediated choice environments.

10. Conclusion

This study finds that AI-enabled smart recommendations are associated with both consumer satisfaction and the structure of digital choice. Personalisation accuracy, transparency, contextual relevance, and diversity each contribute to selected consumer outcomes, while trust and perceived autonomy help explain how recommendation systems influence experience and decision environments. Privacy risk can weaken the satisfaction value of personalisation, whereas digital literacy can strengthen the trust benefits of transparency.

The findings support a balanced view of AI recommendation systems: their value depends not only on predictive accuracy but also on how recommendations are explained, how much variety is offered, how consumer data are handled, and whether users retain meaningful control over the decision environment. The ARSCAM model provides a consolidated framework for examining these relationships in the Indian digital market and offers a basis for future research across platforms, cultures, and emerging generative-AI recommendation systems.

References

- Cai, X., Cebrian, M., & Pentland, A. (2022). Algorithmic defaults and the architecture of digital consumption. *Nature Human Behaviour*, 6(4), 482–491.
- Chung, M., Ko, E., Joung, H., & Kim, S. J. (2020). Chatbot e-service and customer satisfaction regarding luxury brands. *Journal of Business Research*, 117, 587–595.
- Ekstrand, M. D., Harper, F. M., Joshaghani, N., & Konstan, J. A. (2014). User perception of differences in recommender algorithms. *Proceedings of RecSys '14*, 161–168.
- Gedikli, F., Jannach, D., & Ge, M. (2014). How should I explain? A comparison of different explanation types for recommender systems. *International Journal of Human-Computer Studies*, 72(4), 367–382.
- Gomez-Uribe, C. A., & Hunt, N. (2016). The Netflix recommender system: Algorithms, business value, and innovation. *ACM Transactions on Management Information Systems*, 6(4), 1–19.
- Gray, C. M., Kou, Y., Battles, B., Hoggatt, J., & Toombs, A. L. (2018). The dark (patterns) side of UX design. *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*, Paper 534.
- Hair, J. F., Risher, J. J., Sarstedt, M., & Ringle, C. M. (2019). When to use and how to report the results of PLS-SEM. *European Business Review*, 31(1), 2–24.
- Hair, J. F., Hult, G. T. M., Ringle, C. M., & Sarstedt, M. (2021). *A primer on partial least squares structural equation modeling* (3rd ed.). SAGE.
- Henseler, J., Ringle, C. M., & Sarstedt, M. (2015). A new criterion for assessing discriminant validity in variance-based structural equation modeling. *Journal of the Academy of Marketing Science*, 43(1), 115–135.
- Kotkov, D., Wang, S., & Veijalainen, J. (2016). A survey of serendipity in recommender systems. *Knowledge-Based Systems*, 111, 180–192.
- Lee, J., & Kwon, K. N. (2021). Consumer responses to AI-enabled personalized recommendations. *Journal of Retailing*, 97(1), 85–99.
- Logg, J. M., Minson, J. A., & Moore, D. A. (2019). Algorithm appreciation: People prefer algorithmic to human judgment. *Organizational Behavior and Human Decision Processes*, 151, 90–103.
- Longoni, C., Bonezzi, A., & Morewedge, C. K. (2019). Resistance to medical advice from algorithm-driven sources. *Journal of Consumer Research*, 46(4), 629–650.
- MacKenzie, I., Meyer, C., & Noble, S. (2013). How retailers can keep up with consumers. *McKinsey Quarterly*, 4, 1–10.
- Martin, K. D., Borah, A., & Palmatier, R. W. (2017). Data privacy: Effects on customer and firm performance. *Journal of Marketing*, 81(1), 36–58.
- Mathur, A., et al. (2019). Dark patterns at scale: Findings from a crawl of 11K shopping websites. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW), Article 81.
- Mehrotra, R., McInerney, J., Bouchard, H., Lalmas, M., & Diaz, F. (2020). Towards a fair marketplace: Counterfactual evaluation of relevance, fairness and satisfaction in recommendation systems. *Proceedings of CIKM*, 2243–2251.
- Peters, D., Calvo, R. A., & Ryan, R. M. (2018). Designing for motivation, engagement and wellbeing in digital experience. *Frontiers in Psychology*, 9, Article 797.
- Sunstein, C. R. (2020). *Sludge: What stops us from getting things done and what to do about it*. MIT Press.
- Yeomans, M., Shah, A., Mullainathan, S., & Kleinberg, J. (2019). Making sense of recommendations. *Journal of Behavioral Decision Making*, 32(4), 403–414.
- Zhang, Y., Séaghdha, D. O., Quercia, D., & Jambor, T. (2012). Auralist: Introducing serendipity into music recommendation. *WSDM '12*, 13–22.