



Submitted: 20-05-2026 | Revised: 05-06-2026 | Accepted: 10-06-2026 | Published: 30-06-2026

The Privacy Paradox in AI-Driven Marketing: A Theoretical Synthesis of Consumer Trust and Hyper-Personalization

Dr. Sandeep Kumar

Symbiosis Institute of Technology,
Nagpur Campus, Symbiosis International (Deemed University), Pune, India
City- Nagpur, Country-India, Pincode-440008
Orcid: <https://orcid.org/0000-0002-4752-7884>
E-Mail: er.sandeepsahratia@gmail.com

Abstract

Consumers have a high level of privacy concerns and share personal information despite this. This mismatch, the privacy paradox, has been explained by rational tradeoff, bounded rationality, resignation and measurement artefact and continues to be the organizing puzzle in privacy research in marketing. This paper suggests that the puzzle is now mis-specified. All of the most dominant explanations place consumer agency at disclosure, as this was the time when disclosure was the means by which a firm learned about the consumer. Artificial intelligence has broken that connection: modern marketing systems can deduce attributes that have never been revealed, including behavioural data which contains no apparent information. Inference is when a consumer's disclosure behaviour is not contradictory to his/her privacy concern, but the disclosure is merely irrelevant to the outcome that he/she is interested in. This paradox is at least a relic of a disclosure theory in an inference technology. In this paper, the authors introduce the Inferential Trust Model, combining elements of privacy calculus, contextual integrity, and personalization research through three constructs: inferential distance, contextual congruence, and inferential awareness and a mechanism to differentiate between latent and felt privacy violations. The model helps to answer why the paradox is stable (violations are not revealed when the awareness is low), why trust breaks rather than slowly, and why transparency is double-edged. It also suggests an accuracy paradox: The more accurate the model, the harder it is to violate, and the more it is violated, the less trusting the personalization will be. There are then seven propositions and an operationalisation protocol. The report does not include any empirical data, but is, instead, conceptual.

Keywords: *privacy paradox; hyper-personalization; consumer trust; artificial intelligence; contextual integrity; inferential privacy*

1. Introduction

If there's an embarrassing finding in consumer privacy research, it has to be the privacy paradox, which is the most persistent in this field. Consumers believe privacy is important, but they give it up for a discount, convenience or nothing else. This discrepancy was brought up based on the evidence that disclosure intentions and disclosure behaviours are observed as being quite different (Norberg et



al., 2007), and this discrepancy has withstood the scrutiny of two decades of investigation from various disciplines (Smith et al., 2011; Kokolakis, 2017; Barth & de Jong, 2017).

Explanations abound. The privacy paradox does not exist because consumers are rational and make trade-offs between risk and benefit (Dinev & Hart, 2006; Culnan & Armstrong, 1999). The behavioural economics literature suggests that concern and revealed behaviour do not always match, as privacy decisions are context-dependent, malleable, and often uncertain or are influenced by bounded rationality (Acquisti et al., 2015, 2020). Some people think that consumers have just stopped buying. A legal critic argues that the paradox is a chimera: that is, one might compare attitudes towards privacy in general with actions in one particular transaction, and come up with “paradoxes” in any area (Solove, 2021).

All explanations go in the same direction. Each makes the disclosure decision of the consumer the object of the explanation, and asks, “Why didn’t consumers make the disclosure decision based upon stated concern? This was a correct object at the time the theory was developed. In this form, cookie, loyalty card world, what a firm knew about a consumer amounted to little more than what the consumer told it, for the choice whether to disclose meant the difference between privacy and no privacy.

Now it’s not that world anymore. Firstly, AI marketing systems don’t learn from what consumers reveal, but rather what they infer. Such inferences of sensitive personal attributes have been around for more than a decade (Kosinski et al., 2013) and psychological targeting, which exploits such inferences, has clear persuasive power (Matz et al., 2017). Contemporary systems extend this: embeddings, lookalike modelling and behavioural sequence models would identify attribute estimates which no consumer disclosed, would consider as disclosable, and which a consumer could not withhold. The firm’s knowledge and the disclosure of the consumer has parted ways.

This has an impact that the literature has not absorbed. If the firm can deduce from the consumer’s refusal to disclose, then the disclosure decision does not affect the outcome of the privacy and the consumer who discloses but claims that he values his privacy is not being inconsistent. They’re right when they say that if you don’t give, you don’t get. Once the object of analysis shifts from disclosure to inference, the seemingly paradoxical behaviour is rational, or at least unremarkable.

The paper thus poses three questions. RQ1: What happens if disclosure and firm knowledge are separated via inference? RQ2: How does hyper-personalization negatively affect consumer trust when most of the inferences are not visible to the consumer? RQ3: What are the double-sided nature of transparency and when does revealing an inference aid or impede?

It has three contributions: re-specifying the paradox as an artefact of disclosure-centric theory; the Inferential Trust Model which integrates privacy calculus, contextual integrity, and personalization studies into three constructs and a latent–felt violation mechanism; and an accuracy paradox that suggests privacy violation is a mechanism not a remedy (that is, that the accuracy of the model is the way in which privacy is violated rather than how it is achieved). It is a conceptual paper with no data reported.

2. Literature Review

2.1 The Paradox and Its Explanations

Table 1 sets out the dominant explanations and what each assumes.

Table 1. *Explanations of the privacy paradox and their assumptions*

Explanation	Core claim	Locates agency at	Assumes about firm knowledge	Limitation under AI
Privacy calculus	Consumers trade risk against benefit	The disclosure decision	Firm knows what was disclosed	Nothing is traded if nothing can be



Explanation	Core claim	Locates agency at	Assumes about firm knowledge	Limitation under AI
				withheld
Bounded rationality	Preferences are uncertain, context-dependent, malleable	The disclosure decision	As above	Explains inconsistency, not irrelevance
Resignation / cynicism	Consumers believe control is futile	The disclosure decision	As above	Treats futility as belief; under inference it is fact
Measurement artefact	General attitudes vs specific behaviours	The comparison, not the consumer	Agnostic	Correct but incomplete; does not say what to measure instead
Contextual integrity	Violation is flow against context norms	The information flow	Flow is observable	Inference flows are not observable to the consumer

The privacy calculus tradition has conceptualized disclosure as a trade-off between benefits and perceived risks, which is moderated by trust and procedural fairness (Culnan & Armstrong, 1999; Dinev & Hart, 2006), and has provided the industry's most widely used measurement tools (Malhotra et al., 2004). Privacy preferences have been found to be unstable and susceptible to design and framing, as outlined in the behavioural tradition (Acquisti et al., 2015, 2020). The critical tradition claims that the paradox is mainly due to an incommensurability of things compared (Solove, 2021), which this paper endorses and generalizes: the comparison is mistaken because the thing is mistaken.

2.2 Personalization, Trust, and the Creepiness Threshold

Personalization, it has been discovered by marketing research, doesn't always work out in a monotonous manner. Under certain circumstances, personalized advertising yields better results, while in other cases, it has a negative impact: Under conditions of covert data collection (where the user is not aware of the data collection), the negative impact of the personalization paradox (Aguirre et al., 2015) prevails, which leads to a decrease in effectiveness; under conditions of trust building, the negative impact of the personalization paradox is mitigated, while the positive impact of personalized advertising still exists. The key to effectiveness is the match between the personalization, when and where (Bleier & Eisenbeiss, 2015). The impact of privacy on customer and firm performance is measurable (Martin et al., 2017), and privacy is now considered a fundamental marketing issue, not a compliance issue (Martin & Murphy, 2017). Reactivated agency (Tucker, 2014) is a major benefit of privacy controls, as it improves advertising effectiveness.

First, in terms of advertising, transparency reporting the inference flow to users of ads helps or hurts based on the nature of the inference: presenting an acceptable flow of information increases the effectiveness of the ad, whereas revealing an unacceptable flow of information decreases the effectiveness of the ad (Kim et al., 2019). This discovery is typically interpreted as a design lesson. Herein lies the clue, argues this paper: Transparency does not alter the inference, just its awareness; awareness has an effect, so that's the whole story.

2.3 AI in Marketing and the Emergence of Inference



Marketing scholarship has explored strategic and experiential aspects of AI and what it means for marketing tasks, how it will affect the customer experience, and how it is likely to transform marketing practice in general. Marketing scholarship has examined strategic and experiential implications, AI's impact on marketing tasks and its potential for reorganisation of the customer experience, and its likely impact on marketing practice as a whole (Huang & Rust, 2021; Puntoni et al., 2021; Davenport et al., 2020). The closest to the present argument, according to Puntoni et al. (2021), are activities of data capture and data classification that are experiential facts for consumers, an area in which the privacy paradox literature has yet to touch.

2.4 Contextual Integrity

The disclosure frame's most theoretically fruitful alternative is that privacy is not about hiding or controlling information, but rather the rightness of information flows in terms of the norms of the context it enters. If it is a medical fact shared with a doctor, it is okay if it is shared with an advertiser, even with consent it is not. Contextual integrity moves the analysis away from the individual's decision and into the legitimacy of the flow: this paper takes steps towards doing just that. Its constraint is that it assumes that flows can be identified. An inferred attribute is one the consumer isn't able to see that the firm has deduced.

2.5 The Gap

The literature on privacy theorises disclosure, the personalization literature theorises reactions to visible personalization, and the literature on AI theorises capability. No one posits the inference as the privacy-relevant action and no one asks what happens to a privacy-relevant paradox over disclosure if disclosure ceases to be the basis for knowledge.

3. Research Methodology

3.1 Design

Conceptual paper with theory synthesis – Synthesis of several disconnected literatures into a new framework not provided by any of the literatures alone (Jaakkola, 2020). The approach is appropriate if constructs are well developed, but their relations are unspecified, as is the case here where privacy calculus, contextual integrity, personalization effects and AI capability are well-developed but unconnected with one another. It judges against the standard criteria for conceptual contribution – constructs (what), relations (how), justification (why) and boundary conditions (when, where, who) (Whetten, 1989) – and by the requirement that the contribution do identifiable work, rather than summarise (MacInnis, 2011).

3.2 Procedure

The development went through 4 phases. The first step was to analyse the explanations of the paradox, revealing a common assumption that is reported in Table 1 and is the diagnostic move in the paper. Second, constructs were selected from relevant literatures and harmonised because there are different referents for the terms “privacy concern,” “vulnerability,” and “violation” between the different literatures. Third, the constructs were combined into a model following the principle that a violation had to be perceived for it to impact trust, a separation that the literature did not make because violations were mostly perceptible by construction under disclosure. Fourth, a protocol for operationalisation was developed for each proposition and a proposition was derived for each relation.

3.3 Scope and Boundary Conditions

The model deals with commercial marketing settings where a company possesses behavioural data and uses inferential models on it. It doesn't explain state surveillance, in which the consent structure and its implications vary, or explicit disclosure of sensitive data, in which classic privacy calculus holds. No



empirical data collected and no estimates reported: quantities in the operationalisation protocol are in placeholder notation.

4. The Inferential Trust Model

4.1 The Decoupling

Structural statement of the paper's foundational claim is given in figure 1. The classic model is the disclosure model in which the consumer reveals, the firm is aware of what they reveal, and the consumer's decision to disclose is the point of control—the less privacy concern, the more disclosure; the more privacy concern, the less disclosure—and the paradox is that. In the inferential model, the client leaves behind behavioural residues when using the service, the service provider infers properties of the service from these residues, and the client's disclosure decision only affects a diminishing portion of the service provider's knowledge.

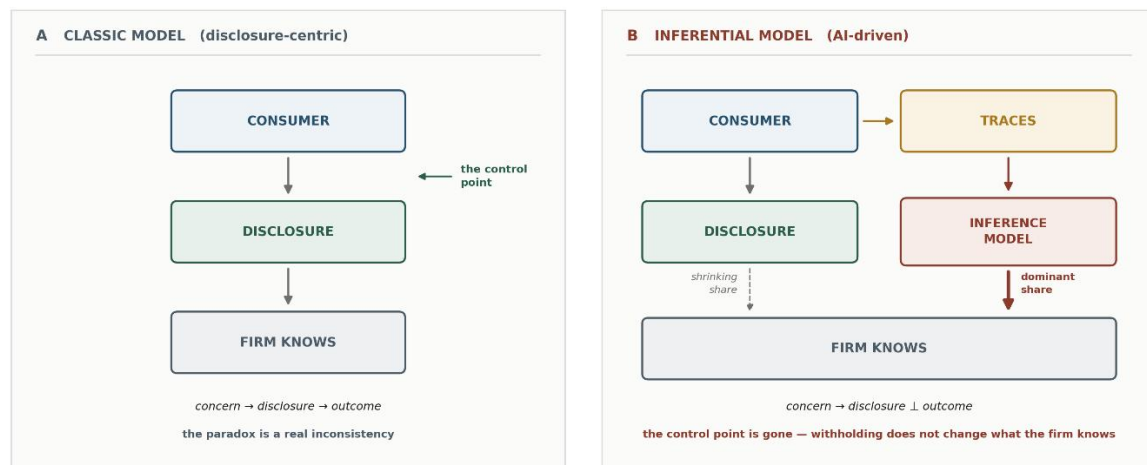


Figure 1. *The decoupling of disclosure from firm knowledge: the classic model and the inferential model.*

The result is the definitive follow. With decoupling, there is no contradiction in “I care about privacy but I disclosed anyway”, as there was no commitment to disclosure. If the consumer does not state their age, it is assumed to indicate the consumer falls within an age band and if the consumer does not state his/her income, the consumer is scored on the income. Withholding is not a strategy, it's a gesture. This brings the paper's first conclusion and recasts the paradox as a product: what is measured is no longer able to affect what is measured and what is measured is a puzzle about consumers.

4.2 Constructs

Table 2. *Constructs of the Inferential Trust Model*


Construct	Definition	Indicative measurement
Inferential distance (ID)	The semantic and statistical distance between what a consumer disclosed	Number of inferential steps from disclosed datum to targeted attribute;



Construct	Definition	Indicative measurement
	and what the firm infers	predictive gain of the model over disclosed data alone
Contextual congruence (CC)	The degree to which an inference flow conforms to the norms of the context in which the source data were generated	Consumer-rated appropriateness of the flow, elicited against the originating context
Inferential awareness (IA)	The extent to which a consumer knows that a given inference is being made and is possible	Elicited estimates of what the firm can infer, benchmarked against what it does infer
Latent violation	An incongruent inference of which the consumer is unaware	Product of low CC and low IA
Felt violation	An incongruent inference the consumer perceives	Product of low CC and high IA
Trust	Willingness to be vulnerable to the firm's use of data	Established trust and vulnerability instruments adapted to the inference context

Latent and felt violation is the model's main tool and follows from the decoupling. In disclosure, a violation was generally obvious, in that the customer was conscious of what they were giving. With inference, the majority of violations are not visible. It is not a violation that is not perceived that harms trust, it is experience, and it is not irrationality of consumers that is the reason for the stability of the paradox.

4.3 The Mechanism



A violation must be perceived before it affects trust. Low awareness keeps violations latent — and the paradox stable.

Figure 2. *The Inferential Trust Model: from inference to trust via latent and felt violation.*

This is the run of the model. Behavioural traces are input to an inferential model which generates an attribute estimate at some inferential distance from the disclosed. The inference is congruent with



norms of originating context or it is not, and in the latter case, it is a violation. Inferential awareness is likely to be low, and firms have all the incentive to keep it low, so that will make a difference or not in terms of trust. If the awareness is low, the violation remains unexposed, the trust remains intact and the consumer tells the paradox happily in its traditional form. The more we're aware, the more it becomes felt and the more trust comes back. Revelation events raise awareness: An ad that is wonderfully targeted, a transparency disclosure, a data breach, a press investigation or a regulatory requirement. Revelation does not cause the violation, it exposes a violation that was already happening. This is why the disclosure frame's 'Consumer Trust in Data-intensive firms' score has a tendency to drop precipitously, rather than gradually. On the day of the scandal, there was no change of conduct within the firm. The difference was awareness and the inventory of lurking offenses was instantly apparent.

4.4 The Accuracy Paradox

The model indicates some unnaturalness for practice. In the conventional industry story, personalization goes awry when it gets a little wrong: When it's not relevant, ads can be a nuisance; when it's better, it's better for consumers; and when it's accurate, it's better for them. The model is the opposite of this. An inference feels invasive because of its inferential distance and accuracy is made from an inferential distance. By construction a model that predicts a sensitive attribute from mundane traces, the more accurate the model, the more it knows that was never given and the more likely a revelation event is to expose an incongruent flow. A high degree of accuracy also leads to a higher probability of revelation: If the inference is very accurate, the consumer is likely to be impressed by the perfectly suitable ad that will make them wonder how the firm knew. The extent to which the ad is badly targeted is forgettable, the perfectly targeted one is a revelation event. Under low contextual congruence, it is therefore the accuracy that is the mechanism of violation, rather than the remedy. Having congruence is not a negative, because when it's high the accuracy is undeniably good.




Figure 3. *The theorised personalization–trust relationship under high and low contextual congruence, with the revelation discontinuity.* Schematic and ordinal: curves represent theorised relationships, not estimated values.

The result is shown in Figure 3: When congruence is high, trust increases as personalization increases; when congruence is low, trust increases as violations grow but it drops off sharply at the point of revelation, and the more the stock of violations grows, the more severe the drop will be. The company which has been most successful in being invisible has the most to lose.

4.5 Propositions



Table 3. *Propositions of the Inferential Trust Model*

ID	Proposition
P1	The association between privacy concern and disclosure behaviour weakens as inferential distance increases; the privacy paradox is strongest where disclosure matters least.
P2	Trust damage is a function of contextual congruence, not of data volume: a small incongruent inference damages trust more than a large congruent one.
P3	Inferential awareness moderates the effect of incongruent inference on trust; where awareness is low, incongruent inference has no trust effect.
P4	Revelation events convert latent violation into felt violation, producing discontinuous rather than gradual trust decline.
P5	The magnitude of trust loss at revelation increases with the accumulated stock of latent violation.
P6	Transparency disclosure raises inferential awareness and is therefore beneficial under high contextual congruence and harmful under low congruence.
P7	Model accuracy increases trust under high congruence and decreases it under low congruence (the accuracy paradox), because accuracy raises both inferential distance and revelation probability.

4.6 Operationalisation

Table 4. *Constructs, measurement approach, and placeholder estimands*

Relation	Design	Estimand (placeholder)
Concern–disclosure attenuation (P1)	Experimental variation of inferential distance; measure concern and disclosure	β_{ID} = [to be estimated]
Congruence vs volume (P2)	Factorial manipulation of flow congruence and data volume	Δ_{CC} vs Δ_{vol} = [to be estimated]
Awareness moderation (P3)	Awareness manipulated via disclosure; incongruent inference held constant	γ_{IA} = [to be estimated]
Revelation discontinuity (P4, P5)	Longitudinal design with an exogenous revelation event	δ_{rev} = [to be estimated]
Transparency conditionality (P6)	Transparency $\times$ congruence factorial	θ_{TC} = [to be estimated]
Accuracy paradox (P7)	Accuracy $\times$ congruence factorial with real inference	λ_{acc} = [to be estimated]

P6 is directly testable with established methods, since transparency effects have been shown to depend on what the explanation reveals (Kim et al., 2019); the model’s contribution is to specify congruence, rather than intuition about “creepiness,” as the moderator.



5. Discussion

5.1 Theoretical Implications

There are three moves in the synthesis. It moves the focus of privacy research from disclosure decision to inference, and turns the 20 years of paradox finding into measures of a decision now made for decoration. It separates violation from perception, disclosure frame could not do, that explains the stability of the paradox, that of continuity of the collapse of trust. It depends on congruence not on volume or accuracy, and that's the two sides of the conditional relationship that helps (Bleier & Eisenbeiss, 2015) as well as that that hurts (Aguirre et al., 2015).

The model also refutes Solove's (2021) criticism. He's correct that the paradox is mostly myth, but it's more than just a mismatch in measurement; it's the behaviour no longer influenced the outcome.

5.2 Practical and Policy implications

The awkwardness of P7 is that the firm optimizing for accuracy at low congruence has a liability that is accruing for them, and that liability will be paid out in full. The defensible strategy is not less accuracy – it is congruence inferences that consumers would back if they were told – a testable criterion, not a slogan.

For policy, the implication is that there is a notice-and-choice regime that applies to the wrong act. Consent is consent to disclosure; disclosure is no longer where privacy is determined; a regime which perfects consent, unperfected disclosure. Regulating inference which attributes may be inferred, in which contexts, regardless of the lawfulness of the source data is the intervention the model implies and is a lot more difficult than what is currently attempted.

S. requirements regarding transparency should be given attention in particular. P6 assumes that enforced transparency makes things more visible and turns previously invisible violations into visible ones. It's not a justification for not being transparent: If there's a violation that's felt, it's a violation that can be contested; and for not being transparent is not a defensible objective. It is an argument that the degree of transparency that can be seen in practice will have a negative impact on trust because practice is not congruent to transparency, and it will have a negative impact on trust whenever practice is not congruent to transparency.

5.3 Limitations

Model is an untested and conceptual model. Its constructs are demanding: inferential distance requires access to model internals that are not shared by firms, and contextual congruence requires eliciting norms that vary across consumers, cultures and contexts, but measurement is inescapably an aggregation of heterogeneous consumer judgement. The latent/felt distinction is analytically pure and practically odd because the consumer who does not yet have a stake in the trust involved in the latent violation is never the first to discover it. The model also assumes that a perceived violation of trust results in a decrease in trust, but not in a sense as the general awareness of inference increases (which the model does not account for), as that might be the most interesting boundary case, and it could simply fade away as public awareness of how inferences are created becomes more widespread. Finally, the boundary conditions drop state surveillance and explicit contexts of sensitive data, where classic privacy calculus still prevails.

5.4 Future Research

The priority test is P1, which adjudicates the paper's central claim: If the concern–disclosure association does not weaken as the inferential distance increases, the artefact argument fails. The most important one is P7, which is testable in a factorial design with actual inferential models in which real people are actually the subjects, rather than vignettes, a design the field has not embraced due to its ethical implications, which the model amplifies. P4 and P5 are disclosure dependent and should be followed where possible as a disclosure mandate or a publicised disclosure. In addition to the propositions, three directions are of interest; one is whether the level of inferential awareness is increasing within the population as a whole, which would push the entire system toward the felt



branch; another is whether consumers can develop coherent norms that cannot be perceived with the senses alone, ones that have never been tested under the conditions of inference; and another is whether congruence can be predicted well enough to be engineered into the model.

6. Conclusion

The question of why consumers who say they care about privacy actually do not has been a mystery in the field for 20 years. This paper has introduced the idea that the question assumes a world that has been destroyed by AI. If firms learn by inference rather than by disclosure, the consumer's disclosure decision no longer affects what the firm knows and behaviour that is inconsistent with the stated concern will simply be indicative of the irrelevance of the disclosure act being measured. The Inferential Trust Model is based on the premise that the analysis should be moved to the inference, and is framed by the three constructs of inferential distance, contextual congruence and inferential awareness, which provide an account in which the paradox is stable because violations are latent, trust collapses suddenly because revelation makes stocks flow, transparency is double edged because it makes the people aware but does not change their behaviour, and accuracy is dangerous because its meaning is to flow rather than to be stored. The claims are made in an easily disprovable manner. If they are true, the central conundrum in the field is not a consumer conundrum. It's a bit of a mystery that's been going on with a theory that has been watching the hands of the consumer while the firm has learned everything without needing the hands of the consumer.

Declarations

Conflict of Interest

The authors declare no conflicts of interest related to this study.

Funding Statement

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Ethical Approval

This study did not involve any human or animal subjects and, therefore, did not require ethical approval.

References

- Acquisti, A., Brandimarte, L., & Loewenstein, G. (2015). Privacy and human behavior in the age of information. *Science*, 347(6221), 509–514. <https://doi.org/10.1126/science.aaa1465>
- Acquisti, A., Brandimarte, L., & Loewenstein, G. (2020). Secrets and likes: The drive for privacy and the difficulty of achieving it in the digital age. *Journal of Consumer Psychology*, 30(4), 736–758. <https://doi.org/10.1002/jcpy.1191>
- Aguirre, E., Mahr, D., Grewal, D., de Ruyter, K., & Wetzels, M. (2015). Unraveling the personalization paradox: The effect of information collection and trust-building strategies on online advertisement effectiveness. *Journal of Retailing*, 91(1), 34–49. <https://doi.org/10.1016/j.jretai.2014.09.005>
- Barth, S., & de Jong, M. D. T. (2017). The privacy paradox Investigating discrepancies between expressed privacy concerns and actual online behavior A systematic literature review. *Telematics and Informatics*, 34(7), 1038–1058. <https://doi.org/10.1016/j.tele.2017.04.013>



- Bleier, A., & Eisenbeiss, M. (2015). Personalized online advertising effectiveness: The interplay of what, when, and where. *Marketing Science*, 34(5), 669–688. <https://doi.org/10.1287/mksc.2015.0930>
- Culnan, M. J., & Armstrong, P. K. (1999). Information privacy concerns, procedural fairness, and impersonal trust: An empirical investigation. *Organization Science*, 10(1), 104–115. <https://doi.org/10.1287/orsc.10.1.104>
- Davenport, T., Guha, A., Grewal, D., & Bressgott, T. (2020). How artificial intelligence will change the future of marketing. *Journal of the Academy of Marketing Science*, 48(1), 24–42. <https://doi.org/10.1007/s11747-019-00696-0>
- Dinev, T., & Hart, P. (2006). An extended privacy calculus model for e-commerce transactions. *Information Systems Research*, 17(1), 61–80. <https://doi.org/10.1287/isre.1060.0080>
- Huang, M.-H., & Rust, R. T. (2021). A strategic framework for artificial intelligence in marketing. *Journal of the Academy of Marketing Science*, 49(1), 30–50. <https://doi.org/10.1007/s11747-020-00749-9>
- Jaakkola, E. (2020). Designing conceptual articles: Four approaches. *AMS Review*, 10(1–2), 18–26. <https://doi.org/10.1007/s13162-020-00161-0>
- Kim, T., Barasz, K., & John, L. K. (2019). Why am I seeing this ad? The effect of ad transparency on ad effectiveness. *Journal of Consumer Research*, 45(5), 906–932. <https://doi.org/10.1093/jcr/ucy039>
- Kokolakis, S. (2017). Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers & Security*, 64, 122–134. <https://doi.org/10.1016/j.cose.2015.07.002>
- Kosinski, M., Stillwell, D., & Graepel, T. (2013). Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences*, 110(15), 5802–5805. <https://doi.org/10.1073/pnas.1218772110>
- MacInnis, D. J. (2011). A framework for conceptual contributions in marketing. *Journal of Marketing*, 75(4), 136–154. <https://doi.org/10.1509/jmkg.75.4.136>
- Malhotra, N. K., Kim, S. S., & Agarwal, J. (2004). Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information Systems Research*, 15(4), 336–355. <https://doi.org/10.1287/isre.1040.0032>
- Martin, K. D., Borah, A., & Palmatier, R. W. (2017). Data privacy: Effects on customer and firm performance. *Journal of Marketing*, 81(1), 36–58. <https://doi.org/10.1509/jm.15.0497>
- Martin, K. D., & Murphy, P. E. (2017). The role of data privacy in marketing. *Journal of the Academy of Marketing Science*, 45(2), 135–155. <https://doi.org/10.1007/s11747-016-0495-4>
- Matz, S. C., Kosinski, M., Nave, G., & Stillwell, D. J. (2017). Psychological targeting as an effective approach to digital mass persuasion. *Proceedings of the National Academy of Sciences*, 114(48), 12714–12719. <https://doi.org/10.1073/pnas.1710966114>
- Nissenbaum, H. (2004). Privacy as contextual integrity. *Washington Law Review*, 79(1), 119–157.



-
- Nissenbaum, H. (2010). *Privacy in context: Technology, policy, and the integrity of social life*. Stanford University Press.
- Norberg, P. A., Horne, D. R., & Horne, D. A. (2007). The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs*, 41(1), 100–126. <https://doi.org/10.1111/j.1745-6606.2006.00070.x>
- Puntoni, S., Reczek, R. W., Giesler, M., & Botti, S. (2021). Consumers and artificial intelligence: An experiential perspective. *Journal of Marketing*, 85(1), 131–151. <https://doi.org/10.1177/0022242920953847>
- Smith, H. J., Dinev, T., & Xu, H. (2011). Information privacy research: An interdisciplinary review. *MIS Quarterly*, 35(4), 989–1015. <https://doi.org/10.2307/41409970>
- Solove, D. J. (2021). The myth of the privacy paradox. *George Washington Law Review*, 89(1), 1–51.
- Tucker, C. E. (2014). Social networks, personalized advertising, and privacy controls. *Journal of Marketing Research*, 51(5), 546–562. <https://doi.org/10.1509/jmr.10.0355>
- Whetten, D. A. (1989). What constitutes a theoretical contribution? *Academy of Management Review*, 14(4), 490–495. <https://doi.org/10.5465/amr.1989.4308371>