



(RESEARCH ARTICLE)



FROM CALCULUS TO SUPPRESSION: RECONCEPTUALISING THE PRIVACY TRADE-OFF IN PERSONALISED ADVERTISING UNDER WEAK ENFORCEMENT

Borono Bassey *

Department of Mass Communication, University of Uyo, Nigeria.

* Corresponding Author

ORCID Details

Borono Bassey: ORCID iD: <https://orcid.org/0009-0001-6290-6970>

International Journal of Science and Research Archive, 2026, 20(03), 143–151

Publication history: Received on 28 July 2026; revised on 02 September 2026; accepted on 05 September 2026

Article DOI: <https://doi.org/10.30574/ijrsra.2026.20.3.1714>

Copyright © 2026 Author(s) retain the copyright of this article. This article is published under the terms of the Creative Commons Attribution License 4.0

ABSTRACT

The prevailing account of how consumers make privacy decisions is the privacy calculus, which models disclosure and engagement as a trade in which the benefits of personalisation are weighed against the expected cost to privacy. This paper argues that the calculus is conditional on an institution that most of its evidence has held constant, namely credible enforcement of data protection. Where enforcement is credible, the privacy cost is bounded, because misuse carries consequences and consumers retain recourse, and the calculus behaves as a compensatory trade in which personalisation benefits can offset concern. Where enforcement is weak, the cost is unbounded and non-negotiable, and the trade does not clear. Under that condition privacy concern ceases to function as a term to be weighed and begins to function as a suppressor, one that increases in personalisation cannot offset. The paper develops the construct of privacy suppression, formalises the returns to personalisation as an enforcement-conditional inverted-U in which the peak shifts lower and the decline steepens as enforcement weakens, and proposes a typology of privacy postures adapted to weak-enforcement markets, in which a resigned posture rather than the pragmatist of established segmentations becomes the majority case. Survey evidence from web users in an emerging Nigerian market, where privacy and platform mistrust were the most intensely held attitudes and the strongest negative correlate of purchase, motivates the account. The paper offers six propositions, a conceptual model, boundary conditions, and a testing agenda. Its contribution is to make the privacy calculus conditional on enforcement, and to specify what replaces it where enforcement fails.

Keywords: Privacy calculus; Personalised advertising; Privacy paradox; Data protection enforcement; Digital resignation; Emerging markets

1 INTRODUCTION

Personalised advertising runs on personal data. The practice matches messages to a consumer's inferred interests by collecting and analysing records of browsing, search, and platform behaviour, and its promised advantage over

untargeted advertising rests entirely on the quality and quantity of the data it can draw on (Boerman et al., 2017). Because the practice is built on data, how consumers weigh the exposure of their data is central to whether it works. The dominant account of that weighing is the privacy calculus, which treats disclosure and engagement as the outcome of a trade: consumers accept some loss of privacy when the expected benefit, including the relevance and convenience that personalisation offers, is judged to outweigh the expected cost (Dinev & Hart, 2006). On this account, privacy concern is one term in a calculation, and a sufficiently large benefit can offset it.

This paper argues that the calculus rests on a condition its evidence has rarely had to examine, because that condition was satisfied wherever the evidence was gathered. The condition is credible enforcement of data protection. A calculation of costs and benefits is only well defined when the costs are bounded, and the privacy cost a consumer faces is bounded by the institutions that constrain what may be done with disclosed data and that provide recourse when those constraints are broken. In markets with enforced data-protection regimes, procedural safeguards, and functioning redress, the downside of disclosure is limited and, in a sense, insurable, so a consumer can trade it against the benefits of personalisation in the manner the calculus describes. Those institutions are not universal. In many emerging digital markets, data-protection law is recent or unenforced, procedural safeguards are absent, and redress is unavailable in practice. Where that is so, the privacy cost is not bounded, and a cost that is unbounded cannot be straightforwardly traded against a marginal gain in relevance.

Survey evidence from web users in Akwa Ibom State, Nigeria, shows the consequence (Bassey, 2026). In that market, privacy concern and platform mistrust were the most intensely held of all the attitudes measured, and they were the strongest negative correlate of purchase, exceeding in magnitude the positive contribution of advertising exposure. The benefits of personalisation did not offset this concern; they were overwhelmed by it. Consumers reported greater willingness to purchase from platforms that stated clear privacy practices, which indicates that the concern was responsive to assurance, yet the assurance available from any single platform was not sufficient to restore the trade. This is not the pattern a compensatory calculus predicts, in which a large enough benefit buys disclosure. It is a pattern in which concern dominates the decision and benefit cannot buy its way past it.

The privacy literature has the materials to describe this but has not assembled them into a conditional account. The privacy calculus itself treats concern as a weighable cost (Dinev & Hart, 2006; Smith et al., 2011). The privacy paradox literature documents that expressed concern often fails to change behaviour, and reads the gap as evidence that concern is weak or overridden rather than that it may, under other conditions, dominate (Barth & de Jong, 2017; Norberg et al., 2007). Work on procedural fairness shows that institutional assurance can neutralise concern (Culnan & Armstrong, 1999), which implies, without stating, that the absence of assurance should intensify it. What the field lacks is an account in which the enforcement environment is a variable that governs whether the calculus operates as a trade at all.

This paper offers that account. It argues that the privacy calculus is conditional on enforcement credibility, introduces the construct of privacy suppression to describe the regime that prevails when enforcement fails, formalises the returns to personalisation as an enforcement-conditional inverted-U, and proposes a typology of privacy postures adapted to weak-enforcement markets. The contribution is threefold. First, the paper makes an established decision model conditional, specifying the institutional circumstance under which it holds and the circumstance under which it breaks. Second, it names and characterises what replaces the trade when it breaks, a suppression regime in which concern dominates rather than trades. Third, it revises a standard segmentation of consumers by privacy attitude to fit markets where resignation, not pragmatism, is the majority posture, and it derives propositions and a testing agenda from all three moves.

The remainder of the paper proceeds as follows. Section 2 examines the privacy calculus and the enforcement assumption embedded in it. Section 3 reviews the personalisation-privacy paradox and the conditions under which it has been studied. Section 4 presents the anomaly that weak enforcement produces. Section 5 develops privacy suppression and the enforcement-conditional inverted-U. Section 6 proposes the typology of privacy postures. Section 7 states the propositions and the conceptual model. Section 8 discusses implications, Section 9 sets out boundary conditions and a testing agenda, and Section 10 concludes.

2 THE PRIVACY CALCULUS AND ITS ENFORCEMENT ASSUMPTION

The privacy calculus is the most influential model of how individuals decide to disclose personal information and engage with data-driven services. It holds that disclosure follows a weighing of anticipated benefits against anticipated privacy costs, so that people disclose when the balance is favourable and withhold when it is not (Dinev & Hart, 2006). The model has organised a large and interdisciplinary body of research, which the most cited review classifies and integrates around the antecedents and outcomes of privacy concern and the contexts that shape them (Smith et al., 2011). Its appeal is that it treats the consumer as a reasoning agent rather than either a privacy maximiser or a heedless discloser, and it accommodates the observation that the same person will guard data in one setting and surrender it in another.

The calculus is a calculation, and a calculation of costs presupposes that the costs can be estimated and are bounded. This presupposition is usually implicit, but it is load-bearing. The privacy cost of disclosure is not intrinsic to the act of disclosing; it is a function of what can subsequently be done with the data and of what recourse the discloser retains if the data are misused. Both are set by institutions external to the individual transaction. Where a data-protection regime constrains secondary use, mandates security, and imposes penalties for breach, the plausible downside of disclosure is limited, and the consumer can weigh a bounded cost against the benefit of personalisation. The clearest demonstration of this dependence is the finding that procedural fairness neutralises privacy concern: when consumers are assured that fair information practices govern the use of their data, concern ceases to distinguish those willing to be profiled from those unwilling (Culnan & Armstrong, 1999). Assurance works by bounding the cost, and bounding the cost is what makes the calculus a trade.

Two features of the tradition follow from this and matter for the argument. The first is the settings in which its evidence was produced. The calculus has been elaborated and tested overwhelmingly in mature markets with enforced privacy regimes, whether the sectoral protections and procedural norms of the United States or, more recently, the comprehensive regimes of Europe. In those settings the bounding institutions are present by default and do not vary within a study, so their role in making the calculus operate is not observed. The second is the treatment of concern. Across the tradition, privacy concern is modelled as a magnitude, a larger or smaller cost that enters the weighing (Smith et al., 2011). Behavioural research complicates this picture by showing that concern is uncertain, context-dependent, and malleable, manipulable by the very firms and governments whose conduct it concerns (Acquisti et al., 2015). Yet even this behavioural turn locates the variation in the individual's psychology, in the uncertainty and constructed preferences of the person, rather than in the institutional environment that determines whether the cost the person contemplates is bounded at all. The step this paper takes is to treat that environment as the variable.

2.1 The Personalisation-Privacy Paradox and the Conditions of Its Study

The most discussed empirical regularity in this area is the privacy paradox, the repeated finding that expressed privacy concern is a poor predictor of disclosure behaviour, so that people who report high concern nonetheless surrender data readily (Norberg et al., 2007). A systematic review of the paradox confirms its persistence across contexts and catalogues the explanations offered for it, from the discounting of distant risks to the salience of immediate rewards (Barth & de Jong, 2017). In the advertising literature the paradox appears as the personalisation paradox, in which the data collection that makes advertising relevant also makes consumers feel exposed, so that covert collection can depress the response that relevance would otherwise produce (Aguirre et al., 2015). Related work shows that the effectiveness of personalised advertising is contingent rather than automatic, rising when consumers are granted control over their information (Tucker, 2014).

Read carefully, this literature contains a condition it does not foreground. The paradox is defined as a gap between high concern and disclosure that occurs anyway, and the explanations for it generally assume that disclosure is the behaviour to be explained, treating concern as the puzzle because it fails to prevent an action that is otherwise expected. That framing makes sense in a setting where disclosure is safe enough to be the default and concern is the anomaly that ought, but fails, to disrupt it. It makes much less sense in a setting where the safety of disclosure cannot be assumed. In such a setting, concern that suppresses engagement is not paradoxical at all; it is the rational response, and the behaviour to be explained is disclosure, not restraint. The paradox literature, in other words, is a literature about markets in which disclosure is the sensible baseline, and its central puzzle presupposes the enforcement conditions this paper treats as variable.

The point is reinforced by the observation that privacy concern is context-dependent and malleable (Acquisti et al., 2015). If concern is shaped by context, then the enforcement environment is one of the most powerful features of context that could shape it, and a body of work conducted almost entirely within one kind of enforcement environment cannot establish how concern behaves across environments. The condition under which the paradox holds, that disclosure is safe enough to proceed despite concern, is precisely the condition that weak enforcement removes.

3 THE ANOMALY UNDER WEAK ENFORCEMENT

The emerging-market evidence does not display the privacy paradox. It displays something closer to its inverse. In the survey of web users in Akwa Ibom State, privacy concern and platform mistrust were not weak predictors overridden by the convenience of disclosure; they were the most intensely held attitudes measured and the strongest negative correlate of purchase (Bassey, 2026). Concern did not fail to change behaviour. It changed behaviour more than the benefits of personalisation did, and in the opposite direction. Consumers exposed to relevant, behaviourally targeted advertising withheld purchase, and the more the environment signalled that their data were being gathered and used, the more that signal worked against response rather than for it.

Two details sharpen the diagnosis. First, concern was responsive to assurance: consumers reported greater willingness to purchase from platforms with clearly stated privacy practices (Bassey, 2026). This rules out the interpretation that concern was simply a fixed disposition immune to institutional signals, and it shows that the suppression was tied to the perceived safety of data handling rather than to a blanket refusal to transact. Second, the assurance that any single platform could offer was not enough to restore the trade. A stated privacy policy raised willingness at the margin but did not convert the intensely concerned into ready disclosers, which is what one would expect if the missing element were not a private promise but a credible public guarantee that such promises would be honoured. The Nigerian data-protection statute, enacted recently and administered by an authority still building its capacity, exists on paper without yet supplying that guarantee in the perception of consumers (Nigeria Data Protection Act, 2023).

This pattern cannot be accommodated by a compensatory calculus. In a trade, a large enough benefit purchases disclosure, and concern that is responsive to assurance should be tradeable against relevance. Here, benefit did not purchase disclosure, and responsiveness to assurance coexisted with a refusal that private assurance could not overcome. The behaviour is consistent instead with a regime in which the privacy cost is not a bounded term to be weighed but an open-ended liability that dominates the decision. Characterising that regime is the task of the next section.

4 PRIVACY SUPPRESSION AND THE ENFORCEMENT-CONDITIONAL INVERTED-U

This paper proposes privacy suppression to describe the regime that prevails when enforcement fails. Privacy suppression is a state in which privacy concern operates as a suppressor of engagement rather than as a cost to be traded against benefit, such that increases in the benefits of personalisation do not offset it and may deepen it. The defining feature is not that concern is large. A large but bounded concern can still be traded, and a sufficiently large benefit can still win. The defining feature is that concern has ceased to be compensable, because the cost it represents has no ceiling that the consumer can rely on. The claim is not that the objective harm is infinite but that, absent enforcement, the downside is uninsurable: there is no bound set by recourse, no assurance that a stated limit will hold, and so no figure a consumer can enter into a trade with any confidence. When the downside cannot be bounded in this sense, no marginal increase in relevance closes the gap, and the calculation that the privacy calculus describes does not resolve in the way a trade resolves.

The mechanism has a direct implication for the returns to personalisation, which the standard account treats as increasing. If greater personalisation both raises perceived relevance and, by demonstrating how much the advertiser knows, raises the salience of surveillance, then its net effect depends on which of these grows faster and on whether the surveillance it reveals is bounded by recourse. Where enforcement is credible, the surveillance signal is tolerable and relevance dominates, so response rises with personalisation, consistent with the meta-analytic finding that personalisation persuades on average and does so through relevance (Yeo et al., 2025). Where enforcement is weak, the surveillance signal is not bounded, and beyond a modest level of personalisation it comes to dominate, so response first rises and then falls. The returns to personalisation are therefore an inverted-U whose shape is conditioned by enforcement: as enforcement credibility declines, the peak shifts toward lower levels of personalisation and the decline beyond it steepens, until under sufficiently weak enforcement the ascending portion nearly disappears and additional personalisation mostly suppresses.

This account assembles findings that the existing literature has reported separately as symptoms of one underlying regime. Reactance to over-personalisation, observed when highly personalised messages are not justified and their utility is low, is the descending portion of the curve in a setting where justification is scarce (White et al., 2008). The backfire of high personalisation once privacy concern is activated is the same descending portion produced experimentally by switching concern on (Kim & Han, 2025). The instability of the relevance benefit when personalisation reads as creepiness is the flattening of the ascending portion (De Keyser et al., 2024). The depression of response by covert data collection is suppression triggered by the discovery that the cost was larger than assumed (Aguirre et al., 2015). Each of these is a local observation of a curve that this paper describes globally and ties to a single moderator. The role of trust in conditioning whether personalisation helps or harms, established for advertiser-level trust (Bleier & Eisenbeiss, 2015), is here raised to the level of the enforcement regime that governs whether concern can be bounded at all.

Privacy suppression should be distinguished from a merely high privacy weight within the calculus. A high weight preserves the structure of the trade; it raises the benefit required to secure disclosure but leaves disclosure purchasable in principle. Suppression alters the structure, because an uncompensable cost is not a weight in a sum but a constraint that the sum cannot overcome. The distinction is testable, since a high weight predicts that sufficiently large benefits restore disclosure while suppression predicts that they do not, and it matters for practice, because a high weight invites firms to offer more benefit while suppression tells them that benefit is not the binding lever.

4.1 A Typology of Privacy Postures Under Weak Enforcement

If enforcement governs whether the calculus operates as a trade, then the distribution of consumer orientations toward privacy should differ systematically between strong and weak enforcement regimes, and the segmentations built in strong regimes should travel poorly. The standard segmentation, derived from decades of survey research, divides consumers into privacy fundamentalists, who are highly protective; privacy pragmatists, who weigh costs and benefits case by case; and the privacy unconcerned, who disclose readily (Kumaraguru & Cranor, 2005). This segmentation is organised around the pragmatist, the case-by-case weigher who is assumed to be the modal consumer, and the pragmatist is precisely the figure the privacy calculus describes. The segmentation therefore inherits the calculus's enforcement assumption, because the pragmatist can only weigh case by case if the cost is bounded well enough to be weighed.

Where enforcement is weak, a posture that these segmentations treat as marginal becomes central. Consumers who hold persistent, intense concern but disclose anyway, not because they have weighed and accepted the cost but because they see no way to avoid it, occupy a position that has been described under several names, including privacy cynicism, an attitude of powerlessness, mistrust, and resignation that renders protection subjectively futile (Hoffmann et al., 2016), and digital resignation, the condition of wanting to control one's information but feeling unable to, which corporate practices actively cultivate (Draper & Turow, 2019). The resigned consumer is not a pragmatist, because no case-by-case weighing is occurring; disclosure is experienced as unavoidable rather than chosen. Nor is the resigned consumer unconcerned, since concern remains high. Resignation is a distinct posture, and the argument of this paper predicts that it is the majority posture where enforcement is weak, because it is the natural adaptation to an uncompensable and inescapable cost.

A typology adapted to weak-enforcement markets therefore retains the fundamentalist and the unconcerned but replaces the central pragmatist with two postures that the breakdown of the calculus produces. The resigned disclose under protest, holding high concern that finds no effective outlet, and their disclosure should not be read as consent to a trade. The defensive verify and self-protect, translating concern into effortful scrutiny and selective withdrawal rather than resigned disclosure, and their behaviour raises the friction of every transaction. The relative sizes of these postures are themselves a function of enforcement credibility, so that as enforcement improves the resigned and defensive give way to the pragmatist, and as it deteriorates the pragmatist gives way to them. The typology is thus not a fixed taxonomy of persons but a distribution that shifts with the institutional environment.

The resigned posture also resolves an apparent tension in the account. If concern suppresses response, why do resigned consumers disclose their data at all? The answer is that disclosure and purchase are different decisions with different degrees of freedom. Participating in digital life, and so disclosing the data that platforms collect as a condition of access, is largely non-discretionary, since the services on which social and economic life increasingly depend cannot be used without it, and privacy cynicism and digital resignation are precisely the coping attitudes that accompany this unavoidable disclosure (Draper & Turow, 2019; Lutz et al., 2020). The discretionary purchase, by contrast, is a decision the consumer can withhold, and it is there that suppression shows itself. Privacy suppression is therefore not a claim that the concerned refuse to disclose; it is a claim that they withhold the discretionary response, the purchase, that the disclosed data were meant to produce. Resignation explains the continued flow of data, and suppression explains the missing conversion, and the two are the same regime seen from the two sides of a decision the consumer is and is not free to refuse.

4.2 Propositions and Conceptual Model

The following propositions state the account's claims.

- **Proposition 1.** The privacy calculus is conditional on enforcement credibility. Where enforcement is credible, disclosure and engagement reflect a compensatory trade of personalisation benefit against a bounded privacy cost; where enforcement is weak, the cost is unbounded and the trade does not clear.
- **Proposition 2.** Under weak enforcement, privacy concern operates as a suppressor of response rather than as a term traded against personalisation benefit, so that increases in benefit do not offset it.
- **Proposition 3.** The returns to personalisation are non-monotonic and enforcement-conditional. Response rises with personalisation up to a threshold and falls beyond it, and as enforcement credibility declines the threshold shifts lower and the decline steepens, until under sufficiently weak enforcement additional personalisation predominantly suppresses response.
- **Proposition 4.** Procedural assurances offered by an individual firm, including privacy notices, consent mechanisms, and transparency, raise response at the margin but cannot fully substitute for absent public enforcement, so their effect is bounded where enforcement is weak.

- **Proposition 5.** The distribution of privacy postures shifts with enforcement credibility. Where enforcement is weak, a resigned posture, in which persistent high concern coexists with disclosure experienced as unavoidable, predominates over the case-by-case pragmatist of established segmentations.
- **Proposition 6.** Privacy suppression is strongest for exchanges that are data-sensitive or high in consequence and weakest where disclosure is minimal or the data involved are non-sensitive.

Figure 1 presents the enforcement-conditional inverted-U at the centre of the account, showing response to personalisation under high, moderate, and weak enforcement, together with the shift of privacy concern from a tradeable term to a suppressor as enforcement declines.




Figure 1: The enforcement-conditional returns to personalisation. As enforcement credibility declines, the peak of the personalisation-response curve shifts to lower levels of personalisation and the decline beyond it steepens; under weak enforcement, privacy concern operates as a suppressor rather than a tradeable cost.

Propositions 1 and 2 state the conditional breakdown of the calculus and the shift of concern from term to suppressor. Proposition 3 gives the breakdown its observable signature in the returns to personalisation. Proposition 4 bounds the remedy available to firms and locates the binding remedy in public enforcement. Propositions 5 and 6 carry the account into the distribution of consumer postures and the scope conditions of suppression.

5 DISCUSSION

5.1 Theoretical Implications

The primary contribution is to make an established decision model conditional. The privacy calculus has functioned as a general account of privacy decision-making, and the personalisation-privacy paradox has functioned as a general regularity, but both are, on the argument here, descriptions of a particular institutional regime rather than universal laws. By identifying enforcement credibility as the condition under which the calculus operates as a trade, the paper specifies where the model holds and where it fails, and it turns the paradox from an anomaly of human psychology into a feature of markets with functioning enforcement. This is a reinterpretation of a large body of evidence rather than a rejection of it. The compensatory trade that the calculus describes is real where its enforcement condition is met, and the suppression regime this paper describes is what the same consumers exhibit where the condition is not met.

The account also gives privacy suppression a precise relationship to adjacent constructs. It is not privacy concern, which is a magnitude the calculus already contains, but a change in how that magnitude enters the decision, from a weighable cost to an uncompensable constraint. It draws on the finding that procedural fairness neutralises concern (Culnan & Armstrong, 1999) and reads its converse, that the absence of enforceable fairness renders concern uncompensable, and it draws on the behavioural insight that concern is context-dependent and malleable (Acquisti et al., 2015) while insisting that the most consequential feature of context is the enforcement environment rather than the framing of a

single choice. It is related to institution-based trust, the belief that structural assurances make an environment safe (McKnight et al., 2002), but it is narrower and differently placed: it concerns specifically the enforcement of data protection, and it enters as the condition that determines whether the privacy cost can be bounded, not as a general antecedent of transaction trust. The typology of postures, finally, contributes a corrective to segmentations that treat the pragmatist as modal, by showing that the modal posture is a function of the enforcement regime and that resignation, not pragmatism, is modal where enforcement is weak (Draper & Turow, 2019; Hoffmann et al., 2016).

Privacy suppression is also distinct from the coping constructs on which the typology draws. Privacy cynicism, privacy fatigue, and digital resignation name individual attitudes or emotional states, the powerlessness, exhaustion, and futility with which people cope with pervasive data collection (Choi et al., 2018; Draper et al., 2024; Hoffmann et al., 2016). Privacy suppression is a claim at a different level. It is not an attitude a person holds but a property of the decision that an enforcement environment produces, in which concern enters as an uncompensable constraint rather than a weighable cost. The coping attitudes are, on this account, the individual-level correlates of the regime, and the resigned and defensive postures are the forms they take, but the suppression itself is a feature of the market's institutions rather than of any consumer's psychology.

This levelling has a critical implication that the disclosure-suppression distinction brings into view. Because resigned consumers continue to disclose even as they withhold purchase, the two things a data-driven advertising business wants from them come apart. The continued flow of data sustains the profiling on which the business is built, while the withheld purchase denies it the conversion that profiling is meant to deliver. Firms therefore hold a mixed and partly self-defeating incentive. The routine practices that cultivate resignation, opaque collection, unreadable notices, and consent that is nominal rather than real, secure the data by making resistance feel futile (Draper & Turow, 2019), but the same futility is what turns concern into a suppressor of the purchase. An industry optimised to extract disclosure may thus be manufacturing the very suppression that caps its advertising returns. The malleability of privacy concern cuts both ways here (Acquisti et al., 2015): concern that can be dampened into resignation can also, through credible assurance, be restored to a tradeable cost, and where the first path secures data at the expense of conversion, only the second reopens the trade.

5.2 Practical Implications

The account reorients data strategy in weak-enforcement markets. If concern is uncompensable and the returns to personalisation turn negative beyond a modest threshold, then the reflex to collect more data and personalise more intensely is self-defeating, and data minimisation becomes a response strategy rather than only a compliance obligation. Collecting and visibly using less data lowers the surveillance signal that drives suppression, and it does so more reliably than adding benefit, which the account predicts cannot buy past an uncompensable cost.

The account also identifies the effective lever and its limits. Because suppression is produced by the absence of a bound on the privacy cost, the remedy is whatever supplies that bound. Procedural fairness and transparency at the firm level help, since they were shown to neutralise concern where they are believed (Culnan & Armstrong, 1999), but Proposition 4 holds that their effect is bounded where the public guarantee behind them is missing, because a promise is only as good as the prospect that it will be enforced. The decisive bound is therefore public: credible enforcement of data protection is market infrastructure for personalised advertising, and its absence caps what any individual firm can achieve through its own assurances. This locates part of the responsibility for advertising effectiveness in weak-enforcement markets outside marketing and inside the building of regulatory credibility, and it implies that a statute without visible enforcement, such as a recently enacted regime whose authority is still establishing its capacity (Nigeria Data Protection Act, 2023), will not by itself lift the suppression, because it is enforcement rather than enactment that bounds the cost.

5.3 Boundary Conditions and an Agenda for Empirical Testing

The account is bounded in several ways. Suppression is expected to be strongest for exchanges that are data-sensitive or high in consequence and weakest where disclosure is minimal or the data non-sensitive, so the same consumer may suppress on a purchase requiring identity and payment details while disclosing freely in a low-stakes interaction. The enforcement condition is treated as perceived rather than objective, which is appropriate for explaining behaviour but leaves open the relationship between perceived and actual enforcement, a relationship that firms and governments can influence and that is therefore itself a target of strategy. And the account concerns the enforcement of data protection specifically, so it is distinct from, though often correlated with, the broader reliability of payment, delivery, and seller accountability.

The propositions are open to test. Enforcement credibility can be measured as a perception with items addressing the likelihood that misuse would meet a response, and its interaction with personalisation benefit can be estimated to test whether benefit offsets concern under strong but not weak enforcement, as Propositions 1 and 2 require. Proposition 3 can be tested experimentally by manipulating personalisation intensity across conditions that vary enforcement

credibility and estimating the resulting response curves, which should show a peak that moves lower and a decline that steepens as credibility falls; a re-analysis of the personalisation literature that codes each study for the enforcement environment of its setting offers a further, and readily available, test of whether the ascending relevance effect is confined to strong-enforcement conditions. Proposition 4 can be tested by crossing firm-level procedural assurance with public enforcement credibility and observing whether assurance restores response fully only when enforcement is credible. Proposition 5 can be tested with segmentation designs that include a resignation measure alongside the standard categories and compare posture distributions across markets that differ in enforcement, and Proposition 6 by varying the sensitivity and consequence of the disclosure required. Cross-market comparison, in which the same instruments are fielded where enforcement regimes differ, offers the strongest test of the central claim, since it varies enforcement at the level at which the construct is defined. Because enforcement credibility covaries across markets with income, culture, digital literacy, and the broader reliability of commerce, such comparisons must guard against confounding, whether by measuring these covariates and adjusting for them, by pairing markets that differ in enforcement but resemble one another otherwise, or by combining cross-market data with within-market experiments that manipulate enforcement signals while holding the population fixed.

A caveat concerns measurement of the dependent construct. Because the account predicts a change in the structure of the decision rather than only its level, tests should be designed to distinguish an uncompensable constraint from a merely large cost, for example by establishing whether escalating benefits restore disclosure, rather than measuring disclosure or concern as single quantities.

6 CONCLUSION

The privacy calculus has told a consistent story about how consumers decide to disclose their data and engage with personalised services: they weigh the benefit against the cost and act on the balance. This paper has argued that the story holds only where an institution stands behind it, namely credible enforcement of data protection, which bounds the cost and makes the weighing possible. Where enforcement is weak, the cost is unbounded, the weighing does not resolve, and privacy concern stops behaving as a term in a trade and starts behaving as a suppressor that no increase in personalisation can offset. Under that regime the returns to personalisation invert, the remedy shifts from offering more benefit to supplying the missing bound, and the consumer who discloses does so in resignation rather than agreement. Naming this regime, privacy suppression, and making the calculus conditional on the enforcement that sustains it, extends a model built in well-regulated markets to the many markets that are not, and it turns the intense, purchase-suppressing privacy concern of those markets from a deviation to be explained away into the predictable result of an institutional absence.

STATEMENTS ON COMPLIANCE WITH ETHICAL STANDARDS

Disclosure of conflict of interest

The authors declare no conflicts of interest regarding this manuscript.

REFERENCES

- [1] Acquisti, A., Brandimarte, L., & Loewenstein, G. (2015). Privacy and human behavior in the age of information. *Science*, 347(6221), 509-514.
- [2] 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.
- [3] 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.
- [4] Bassey, B. (2026). Online behavioural advertisements and purchase decisions of web users in Akwa Ibom State [Unpublished doctoral dissertation]. University of Uyo.
- [5] Bleier, A., & Eisenbeiss, M. (2015). The importance of trust for personalized online advertising. *Journal of Retailing*, 91(3), 390-409.
- [6] Boerman, S. C., Kruikemeier, S., & Zuiderveen Borgesius, F. J. (2017). Online behavioral advertising: A literature review and research agenda. *Journal of Advertising*, 46(3), 363-376.
- [7] Choi, H., Park, J., & Jung, Y. (2018). The role of privacy fatigue in online privacy behavior. *Computers in Human Behavior*, 81, 42-51.

- [8] Culnan, M. J., & Armstrong, P. K. (1999). Information privacy concerns, procedural fairness, and impersonal trust: An empirical investigation. *Organization Science*, 10(1), 104-115.
- [9] De Keyser, F., Buzeta, C., & Lopes, A. I. (2024). The role of well-being in consumer's responses to personalized advertising on social media. *Psychology & Marketing*, 41(6), 1206-1222.
- [10] Dinev, T., & Hart, P. (2006). An extended privacy calculus model for e-commerce transactions. *Information Systems Research*, 17(1), 61-80.
- [11] Draper, N. A., Hoffmann, C. P., Lutz, C., Ranzini, G., & Turow, J. (2024). Privacy resignation, apathy, and cynicism: Introduction to a special theme. *Big Data & Society*, 11(3), Article 20539517241270663.
- [12] Draper, N. A., & Turow, J. (2019). The corporate cultivation of digital resignation. *New Media & Society*, 21(8), 1824-1839.
- [13] Hoffmann, C. P., Lutz, C., & Ranzini, G. (2016). Privacy cynicism: A new approach to the privacy paradox. *Cyberpsychology: Journal of Psychosocial Research on Cyberspace*, 10(4), Article 7.
- [14] Kim, H., & Han, S. (2025). Triggering the personalization backfire effect: The moderating role of situational privacy concern. *Behavioral Sciences*, 15(10), 1323.
- [15] Kumaraguru, P., & Cranor, L. F. (2005). Privacy indexes: A survey of Westin's studies (Report No. CMU-ISRI-5-138). Institute for Software Research International, Carnegie Mellon University.
- [16] Lutz, C., Hoffmann, C. P., & Ranzini, G. (2020). Data capitalism and the user: An exploration of privacy cynicism in Germany. *New Media & Society*, 22(7), 1168-1187.
- [17] McKnight, D. H., Choudhury, V., & Kacmar, C. (2002). Developing and validating trust measures for e-commerce: An integrative typology. *Information Systems Research*, 13(3), 334-359.
- [18] Nigeria Data Protection Act, 2023 (Nigeria).
- [19] 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.
- [20] Smith, H. J., Dinev, T., & Xu, H. (2011). Information privacy research: An interdisciplinary review. *MIS Quarterly*, 35(4), 989-1015.
- [21] Tucker, C. E. (2014). Social networks, personalized advertising, and privacy controls. *Journal of Marketing Research*, 51(5), 546-562.
- [22] White, T. B., Zahay, D. L., Thorbjørnsen, H., & Shavitt, S. (2008). Getting too personal: Reactance to highly personalized email solicitations. *Marketing Letters*, 19(1), 39-50.
- [23] Yeo, T. E. D., Chu, T. H., & Li, Q. (2025). How persuasive is personalized advertising? A meta-analytic review of experimental evidence of the effects of personalization on ad effectiveness. *Journal of Advertising Research*, 65(4), 616-631.