



AMERICAN JOURNAL OF BUSINESS AND MANAGEMENT INNOVATION (AJBMI)

VOLUME 1 ISSUE 1 (2026)



PUBLISHED BY
E-PALLI PUBLISHERS, DELAWARE, USA

Trust Rhythm in Human and AI Collaboration for Digital Marketing in Nonprofit and Service Organizations

Alina Demidova^{1*}

Article Information

Received: July 25, 2026

Accepted: October 01, 2026

Published: October 16, 2026

Keywords

*Artificial Intelligence Disclosure,
Brand Trust, Digital Marketing,
Human-Machine Collaboration,
Trust Rhythm*

ABSTRACT

This paper builds a conceptual model that explains how artificial intelligence changes the temporal structure of trust in a brand and defines the threshold beyond which automation erodes trust instead of supporting it. The study applies framework synthesis to peer-reviewed literature published from 2018 onward, supplemented by four earlier works whose definitions and instruments later studies reuse. Antecedents of trust are coded into four operational domains, namely transparency, authenticity, timeliness and human accountability, and converted into six propositions with explicit falsification conditions. A simulation on synthetic parameters tests the internal consistency of the model. Trust under machine mediation depends on three parameters of communication rhythm: frequency, amplitude and predictability. Artificial intelligence raises frequency and predictability because it removes the dependence of regular communication on staff availability. It lowers amplitude because content produced in volume without organizational substance carries no verifiable information. The simulated trust index reaches an interior maximum, after which further automation pushes it below its unautomated level. Disclosure of machine use is not the decisive variable. The decisive variable is whether the audience can read human authorship behind the message. The paper contributes a matrix for allocating marketing tasks between people and machines, indicators for monitoring communication rhythm, and a disclosure rule that reports the configuration of human control rather than the bare fact of technology use. The model is theoretical, its evidence base is concentrated in commercial settings, and the nonprofit case is therefore a boundary condition awaiting longitudinal and experimental testing.

INTRODUCTION

Artificial intelligence has moved from an experimental tool to a standard layer of marketing infrastructure, as chatbots, predictive analytics, sentiment analysis and programmatic advertising now shape how organizations target and personalize messages (Ozikor *et al.*, 2026). Systems write advertising copy, segment audiences, select the moment of contact and answer the first questions of a customer. The mechanisms that produce trust did not change at the same speed. Trust still forms through signals that an audience reads as coming from a responsible party, and its definition as acceptance of vulnerability on the basis of positive expectations about future conduct holds regardless of who produced the message (Mayer *et al.*, 1995). Review-based work in digital marketing positions such systems as amplifiers of human effort with identifiable constraints rather than as substitutes for it (Demidova, 2025a). A review of empirical research reports that trust in artificial intelligence depends on the form in which a system appears to a person and on the transparency and reliability it displays rather than on capability alone (Glikson & Woolley, 2020). This paper begins with the gap between adoption speed and trust formation. A second mismatch compounds it. Most available studies measure the reaction of a person to a single message, while marketing practice builds relationships out of long sequences of messages. A result obtained on one advertisement says little about trust

accumulated over a year.

Two groups of findings pull organizations in opposite directions. Disclosure of machine authorship worsens the evaluation of an advertisement through perceived credibility, and the evidence comes from prosocial advertising, the genre closest to nonprofit communication (Baek *et al.*, 2024). Disclosure of chatbot status before a purchase reduces conversion, although an undisclosed chatbot performs at the level of a competent employee (Luo *et al.*, 2019). The loss attaches to the label and not to the quality of the output. Silence is not a safe alternative. The present model treats undisclosed use as a separate failure once a third party reveals it, since the damage then extends from the message to the channel through which every other claim reaches the audience; this claim is stated below as a testable proposition. Organizations therefore lose trust in two ways at once, by automating more than the audience accepts and by concealing what has already been automated. The technology is not the source of the problem. The source is the configuration in which it is placed, that is, the distribution of authorship and the visibility of that distribution.

The practical problem follows from this configuration. Organizations raise the volume and speed of marketing communication with artificial intelligence and simultaneously lose the part of the signal that carries trust, because the elements an audience reads as human are automated first. The loss is invisible in operational

¹ Independent Researcher, Ukraine

* Corresponding author's e-mail: alinademidovaedu@gmail.com

reporting, since the number of contacts grows while the quality of the relationship declines. Experimental research has established the direction of the disclosure effect and identified human control as its moderator. Applied research has established how automation sustains the frequency of contact under resource constraints (Avagimyan, 2026d). No study connects these results in time. Consequently, no model describes how artificial intelligence changes each parameter of the trust rhythm separately.

The purpose of this paper is to construct a conceptual model that explains how artificial intelligence can support the building of trust between an organization and its audience instead of destroying it, and to specify the threshold beyond which support turns into destruction. Six objectives serve that purpose: to separate trust in a brand from trust in a machine and from adjacent constructs; to describe the continuum of autonomy and locate on it the threshold of readable human authorship; to reduce the documented antecedents of trust to four domains with an admissible level of autonomy attached to each; to show how artificial intelligence moves the three parameters of the trust rhythm in different directions; to formulate a rule for allocating marketing tasks between people and machines; and to specify monitoring indicators and testable propositions. Four research questions follow. RQ1 asks how artificial intelligence should be integrated so that it supports rather than destroys trust. RQ2 asks which elements of communication must remain under human control. RQ3 asks how the process of building trust changes after such systems enter the marketing function. RQ4 asks why nonprofit organizations require particular caution, given that trust, authenticity and human contact form their resource base rather than an addition to it (Cheng & Wang, 2025).

The paper contributes on three levels. Theoretically, it develops the trust rhythm construct, which treats trust as the product of periodic reproduction of signals, for human and machine collaboration and introduces signal authorship as a parameter. Integratively, it links each domain of trust to an admissible level of autonomy, which converts a general recommendation about oversight into a specific constraint. Practically, it offers a task allocation matrix and a set of monitoring indicators. The reference base combines two groups of sources that rarely appear together. Experimental work in marketing and decision psychology explains the reaction of a person to machine authorship (Dietvorst *et al.*, 2015; Puntoni *et al.*, 2021). Applied methodological work on lead generation, retention and attribution explains how regularity of contact is produced in practice (Demydova, 2025a). The first group is strong on mechanism and weak on operations; the second is strong on operations and rarely asks what happens to trust. The value of combining them lies in that asymmetry. The logic of calibrated reliance serves as a reference point throughout, since excessive trust and insufficient trust both carry a cost (Lee & See, 2004). The paper proceeds from a review of the

literature to the method and the simulation specification, then presents and discusses the model, and closes with conclusions.

LITERATURE REVIEW

Trust in a Brand and Trust in a Machine

Trust is defined in the organizational literature as willingness to accept vulnerability on the basis of positive expectations about future conduct (Mayer *et al.*, 1995). Trust is directed at the future and is extended under incomplete information, and neither property disappears when a machine produces part of the communication. The transfer of this definition to human and machine interaction proceeded along two lines. Research on automation treats reliance as a calibration problem, since excessive and insufficient trust both produce losses (Lee & See, 2004). A review of empirical research organizes the evidence by the form of representation of a system and by its capabilities (Glikson & Woolley, 2020). Both lines describe trust in a system.

Marketing poses a different question. An audience does not evaluate the reliability of a recommendation engine. It evaluates the organization that placed the engine into the communication. Research on consumer experience with artificial intelligence separates that experience into distinct types, none of which reduces to a relationship with a brand (Puntoni *et al.*, 2021). The construct used here is therefore trust in the organization under machine mediation, while trust in the system stays outside the dependent variable. The demarcation matters because a large share of publications titled as studies of trust in artificial intelligence in fact measures acceptance of a tool by its user. Acceptance by a marketer and confidence of an audience are different variables.

Consumer Response to Machine Authorship

Decision research contains two results that appear to conflict. People reject an algorithm after observing it err, even when it outperforms a human forecaster (Dietvorst *et al.*, 2015). People also prefer algorithmic advice in tasks of estimation and prediction (Logg *et al.*, 2019). The type of task, not the quality of the studies, resolves the conflict, as the review of task boundaries below shows. Marketing evidence is more consistent, and the two key disclosure studies were introduced above (Baek *et al.*, 2024; Luo *et al.*, 2019). The state of this stream imposes a constraint. Almost every available study is an experiment with a single-exposure design and an immediate measurement. Such a design captures the reaction of a person to a novel signal and says nothing about an audience receiving disclosed machine-assisted communication for a year. Survey evidence from online advertising points the same way: advertising drove purchase intention directly while trust did not moderate the relationship (Aktar & Gacho Segumpan, 2025), which suggests that point measures of intention capture the reaction to a message rather than the accumulated relationship. The absence of longitudinal designs is the principal weakness of the

field and one reason why a conceptual model is required before further experiments are run.

Human Control as a Moderator

The most useful correction to the disclosure literature comes from work on organizational frontlines. Two experiments show that collaboration between a person and a machine reduces negative consumer responses, but only when the configuration signals human control. Machine authorship without such control produces worse responses where perceived morality of the use is low, while messages produced under human control are evaluated as messages of human origin (Haupt *et al.*, 2025). Rejection of an algorithm also weakens once people may modify its output, even slightly, which shows that control need not be extensive to work (Dietvorst *et al.*, 2018). Both results support the treatment of human authorship as a threshold rather than as a binary attribute. Work on public relations and advertising practice reaches a compatible conclusion, since over-reliance on automated systems is reported to erode authenticity and creativity (Ozikor *et al.*, 2026).

Task Boundaries and the Limit of Personalization

Aversion depends on the nature of the task. Resistance is stronger in subjective tasks than in tasks resting on measurable facts (Castelo *et al.*, 2019). Resistance grows sharply when a task touches the identity of a consumer, since automation there is read as a denial of the personal character of the choice (Leung *et al.*, 2018). It also grows where the recipient believes that individual circumstances will be ignored (Longoni *et al.*, 2019). Personalization has a similar limit. Personalized advertising works only where trust in the collector of data is present and turns into a perception of surveillance where it is absent (Aguirre *et al.*, 2015). The same technology therefore strengthens a relationship in one configuration and damages it in another. These four studies form the basis of the allocation matrix presented in the Results and Discussion. Combined, they stop being a list of effects and start behaving like a rule.

The Operational Layer and the Nonprofit Case

Trust is produced by regular communication, and regular communication is produced by an operational system. Sequential deployment of chatbots, automated funnels and application collection systems reports gains at the entry stage of a funnel and in mid-funnel progression (Demidova, 2025b). The path from a content strategy to a completed application is treated as a designed sequence rather than as a set of independent publications (Demidova, 2026). Staged warming outperforms single-exposure contact on advertising platforms, which supports the claim that distributed regular communication beats concentrated intensity (Demidova, 2025a). Regularity usually fails for resource reasons. Predictive scoring architectures for capacity-constrained allocation describe this failure directly (Avagimyan, 2026d), and scoring

models for lead qualification manage it at the level of a single contact (Andreasyan, 2026a). Automation therefore has a legitimate and narrow function, namely removal of the dependence of the communication calendar on staff availability.

Measurement forms the second block. Integration of multi-touch attribution with incrementality testing establishes when a budget decision can be called evidence-based (Andreasyan, 2026d). An attribution protocol for advisory services adapts the same logic to delayed results (Avagimyan, 2026b), and the trace map method converts attribution into client reporting standards (Balaian, 2026a). Diagnostic frameworks linked to key performance indicators supply the structure into which such reporting fits (Andreasyan, 2026b). These sources serve one purpose here: construction of indicators for monitoring the rhythm of communication.

Retention research supplies the third block and a caution. Lifecycle messaging and churn prevention describe the vocabulary of frequency (Andreasyan, 2026e), and a key performance indicator system built around trust, outcomes and renewal shows how retention is measured in advisory relationships (Avagimyan, 2026e). Gamified engagement raises client lifetime value through extrinsic reward (Demidova, 2025b), yet such mechanics operate on loyalty and satisfaction rather than on expectation of future conformity to a mission. Repeat behavior can therefore grow while trust stands still. When repeat purchase is the only proxy, the indicator profile misleads. Channel architecture and positioning form the fourth block. Omnichannel communication with evidence-based segmentation describes the environment in which a rhythm has to be maintained (Avagimyan, 2026c). Trust signals and funnel design are treated as engineered elements of conversion (Balaian, 2026d). Chain models linking a value proposition to a revenue model explain why positioning acquires force only through consistent reproduction (Andreasyan, 2026c; Balaian, 2026c). Diagnostic toolkits for growth bottlenecks name the most common managerial error, namely investment in the easiest domain rather than in the constraining one (Avagimyan, 2026a; Balaian, 2026b).

The nonprofit case sharpens every element above. Research on chatbots in nonprofit organizations examines the antecedents and outcomes of chatbot trust together with social media engagement (Cheng & Wang, 2025). The reading adopted in this paper goes one step further, since the technology is treated here as entering an environment in which trust is the resource base rather than an addition to it. Three structural features distinguish that environment. The resource base rests on voluntary contribution, the result is poorly verifiable by the contributor, and the payer and the beneficiary are different parties.

The review thus locates the gap precisely. Experimental research explains the direction of the disclosure effect and its moderator, and applied research explains how automation sustains frequency. The literature lacks

a model that describes the influence of artificial intelligence on each parameter of the trust rhythm separately and specifies the point at which a gain in frequency is cancelled by a loss of substance in the signal.

MATERIALS AND METHODS

Research Design

The study is conceptual, and the choice follows from the state of the field. Controlled experiments have already estimated the effect of disclosure, so an additional estimate would add a coefficient to a dense set. What the field lacks is an architecture specifying how those coefficients relate over time. The deficit is theoretical and is addressed by theoretical means.

The model was built through framework synthesis in four steps. An initial scaffold was derived from the trust rhythm construct, which treats trust as the product of periodic reproduction of signals rather than as a sum of attributes. Evidence from the reviewed corpus was coded against that scaffold. Codes that did not fit were retained rather than discarded, and the scaffold was revised to absorb them; this is how signal authorship entered the model. The revised structure was converted into propositions stated so that each can be contradicted by observation. The stepwise logic follows the discipline of algorithmic frameworks for marketing audit and prioritization, where every step must be reproducible by another analyst (Balayan, 2026). Framework synthesis was preferred to meta analysis because the reviewed studies measure trust through incompatible instruments, and review based work reports the same heterogeneity across methods and settings (Demidova, 2025a). A fifth step followed the formulation of the propositions. The relationships stated in the model were expressed as functions and evaluated numerically on synthetic parameters to test whether the propositions can hold simultaneously. The simulation plays the role that a formal proof plays in an analytical paper: it demonstrates that the stated relationships are mutually consistent and yield a prediction that observation can contradict. The unit of analysis is the organization as a collective author of communication.

Source Selection and Coding Protocol

Records were retrieved from journals indexed in Scopus and Web of Science through database queries and citation tracking. The window opens in 2018, since the experimental study of consumer response to machine authorship formed a distinct stream from that year onward. Four earlier works were admitted because later studies reuse their definitions or instruments without revision: the integrative model of organizational trust (Mayer *et al.*, 1995), the calibration logic of reliance on automation (Lee & See, 2004), the personalization paradox (Aguirre *et al.*, 2015) and the baseline demonstration of algorithm aversion (Dietvorst *et al.*, 2015). Peer reviewed articles and academic

book chapters in English were retained. Conference abstracts, unreviewed preprints, non indexed outlets and descriptive case reports without a theoretical frame were excluded.

Six search strands structured the retrieval: definitions of trust and their transfer to machine mediated interaction; consumer response to disclosure; human control and shared authorship; task dependence of aversion; operational methods of lead generation, retention and attribution; and adoption of artificial intelligence in nonprofit communication. Two strands returned thin results. Empirical work on artificial intelligence in nonprofit marketing remains scarce, and longitudinal measurement of brand trust under automation is effectively absent. Both gaps are carried into the discussion rather than filled by inference.

Coding proceeded in two passes. The first extracted every reported antecedent of trust with the direction of the effect, the design and the sample. The second reduced the antecedents to domains under a single rule: two antecedents belong to the same domain when they reach the audience through the same class of observable signal. The reduction produced four domains, and no antecedent remained unassigned. Seven comparison vectors were then applied to each item, following the criteria that structure Table 3, and diagnostic frameworks that link criteria to key performance indicators informed their wording (Andreasyan, 2026b). One restriction is deliberate. Industry reports on adoption were read during preparation but were not admitted into the coded corpus, because their sampling procedures are rarely documented. The coded corpus contains thirty five records. Fifteen belong to the theoretical and experimental stream: ten experiments, one survey based study, two reviews and two conceptual works that supply definitions. Twenty belong to the applied methodological stream and describe lead generation, retention, attribution and diagnostic practice. Thirty one records fall inside the publication window, and four earlier works were admitted under the rule stated above. The retrieval was purposive rather than exhaustive, so the corpus supports a conceptual synthesis and does not claim the coverage of a systematic review.

Simulation Specification

Let a denote the share of communication acts published without a human review point, so that a equals zero under full human authorship and one under full autonomy. Three functions translate a into the parameters of the rhythm. Frequency F , measured in communication acts per month, rises linearly from four to twelve acts, because automation removes the dependence of the calendar on staff availability (Equation 1). Amplitude A , the informational weight of a single act, falls quadratically: slowly at first and then rapidly, because the decline accelerates once routine content displaces content with organizational substance (Equation 2). Predictability P rises linearly,

because scheduling reduces the variance of the interval (Equation 3).

$$F(a) = 4 + 8a \quad (1)$$

$$A(a) = 1 - 0.9a^2 \quad (2)$$

$$P(a) = 0.6 + 0.4a \quad (3)$$

The trust increment T accrues multiplicatively rather than additively, because the domains constrain one another instead of substituting for one another, and the logarithm expresses a diminishing return to frequency (Equation 4). The trust index I rescales the increment so that full human authorship equals one hundred (Equation 5).

$$T(a) = \ln \ln [1 + F(a)] \cdot A(a) \cdot P(a) \quad (4)$$

$$I(a) = 100 \cdot \frac{T(a)}{T(0)}, \quad a \in [0, 1] \quad (5)$$

Every functional form and parameter value above is an assumption open to revision. The simulation uses no empirical data, observes no organization and makes no descriptive claim about any real population. It establishes internal consistency and produces a shape that observation can refute.

Limitations of the Method

Four limitations bound the claims that follow. The model rests on secondary evidence and generates no primary measurement, so the propositions are stated but not tested. The corpus concentrates in commercial settings within high income markets, and evidence on

nonprofit communication is limited to a small number of studies (Cheng & Wang, 2025); conclusions about that sector are therefore boundary conditions rather than results. The coding rule that assigns antecedents by class of signal is defensible but not unique, and a rule organized by audience would produce a different partition. Several mechanisms were established far from marketing, in particular in medical decision making, where resistance to machine judgment is strongest (Longoni *et al.*, 2019). The direction of these effects is expected to hold in marketing while their magnitude is not, since the stakes of a purchase differ from the stakes of a diagnosis.

RESULTS AND DISCUSSION

Demarcation of the Construct and the Continuum of Autonomy

Four objects of trust operate simultaneously in machine mediated marketing, and their separation is the first result of this study. Trust in the brand is the expectation that the organization will act in future accordance with its declared mission. Trust in the system is the assessment of accuracy by the person who uses a tool. Trust of the marketer in the tool governs the internal decision to delegate. Perceived acceptability of machine use is the judgment of an audience about whether the technology belongs in this communication, and it moderates every other relationship. Table 1 sets out the holder of each assessment, its evidence base and its status in the model.

Table 1: Objects of trust in machine mediated marketing communication

Object of trust	Holder of the assessment	Evidence base of the assessment	Status in the model
Trust in the brand	Audience of the organization	Consistency between declared mission and observable conduct	Dependent variable
Trust in the system	User of the system	Accuracy and stability of output	Context outside the model
Trust of the marketer in the tool	Marketing team	Result of internal use	Internal loop that shapes configuration
Perceived acceptability of machine use	Audience of the organization	Correspondence between the use and the mission	Moderator of all relationships

The separation matters because applied literature merges these objects and then draws conclusions about one from evidence about another. The second result concerns the form of the automation variable. Autonomy is not binary, and four states can be distinguished. In the assisted state a person writes and the system suggests. In the augmented state the system drafts and a person revises. In the delegated state the system produces and publishes within rules set by a person. In the autonomous state no review point remains. The response of an audience does not fall evenly across these states. It falls sharply where human authorship stops being readable, which lies between the delegated and the autonomous state, as the dependence of consumer response on signalled control indicates (Haupt *et al.*, 2025). A parallel result obtained on decision

makers rather than on audiences points in the same direction, since acceptance of an algorithm returns once a person may modify its output (Dietvorst *et al.*, 2018).

A consequence follows for the disclosure debate. Whether to disclose is secondary. What the disclosure reports is primary. A statement that a machine generated the content gives the audience the label that damages evaluation. A statement describing who set the rules and who checked the result gives the audience the position of the person in the process, which the evidence identifies as decisive.

The Four Domain Architecture

The coding procedure reduced the documented antecedents to four domains. Transparency covers

what an audience can learn about the operation of the organization and about the way its communication is produced. Authenticity covers the sense that a recognizable party stands behind the message. Timeliness covers maintenance of contact at an expected interval and speed of reaction to events; its signal is the interval

rather than the content of any single message. Human accountability covers acceptance of consequences, and its signal is a named person available where a decision is questioned. Table 2 pairs each determinant with its observable signal, its admissible level of autonomy and a proxy indicator.

Table 2: Trust domains, signals, admissible autonomy and proxy indicators

Domain	Determinant	Observable signal	Admissible autonomy	Proxy indicator
Transparency	Configuration of machine use	Statement of who set the rules and who checked the output	Delegated	Share of messages with stated authorship
Transparency	Handling of personal data	Explanation of collection and purpose	Augmented	Rate of consent withdrawal
Authenticity	Voice of the organization	Recognizable style across channels	Augmented	Response to non promotional content
Authenticity	Reference to real cases	Verifiable facts inside the message	Assisted	Share of content with external evidence
Timeliness	Cadence of contact	Interval between publications	Autonomous	Variance of the interval
Timeliness	Reaction to events	Latency of the first response	Delegated	Median response latency
Human accountability	Named responsibility	Signature of a person under a decision	Assisted	Share of escalations closed by a person
Human accountability	Handling of criticism	Reply to a complaint	Assisted	Share of complaints answered by a person

The domains are not substitutes. Transparency without authenticity produces an organization that discloses a great deal and sounds like nobody. Timeliness without human accountability produces a fast reaction for which no one answers. Authenticity without timeliness produces a convincing voice appearing at unpredictable moments, which builds no expectation. The fourth column of Table 2 carries the analytical weight of the study, since it converts a general recommendation about oversight into a constraint attached to a specific signal. The admissible level differs across rows for one reason. Scheduling a publication carries no relational content and therefore tolerates autonomous operation, while a reply to a complaint carries the whole relationship and tolerates none.

Figure 1 assembles these elements into a single model. The left block contains the inputs, that is, the functions in which artificial intelligence operates: analytics, content generation, personalization and contact automation. The centre contains the four domains arranged as a closed cycle rather than as a list, and the three rhythm parameters of frequency, amplitude and predictability sit directly beneath it. The human control loop passes above the centre and marks the signals that cannot bypass a person. The right block contains the outcome, trust in the brand, with three behavioral proxies: repeat support, willingness to share data, and referral and advocacy. Moderators run along the base, namely the objectivity of the task, the relational weight of the act, the sector context and digital literacy, and a

feedback arrow returns from the outcome to the inputs over time. The schematic is necessary because it shows the human loop as a constraint on the domains rather than as an additional input. This distinction separates the present model from accounts that treat oversight as one more resource (Fig. 1).




Figure 1: The Trust Rhythm model of human and machine collaboration in digital marketing

Rhythm Parameters under Automation

The central result concerns the direction in which artificial intelligence moves each parameter. Frequency rises, because automated funnels and application collection systems remove the dependence of the calendar on staff availability (Demidova, 2025b), and staged sequences outperform single exposures precisely because they are sustained (Demidova,

2025a). Predictability rises for the same reason, since a scheduled process misses fewer intervals than a manual one (Avagimyan, 2026d). Amplitude falls. Amplitude is the informational weight of a single act, and content produced in volume without organizational substance carries no verifiable information. Analyses of artificial intelligence in public relations and advertising warn in similar terms that over-reliance on automated systems erodes authenticity and creativity (Ozikor *et al.*, 2026). A message that reports nothing new occupies an interval

without filling it. The organization keeps the calendar and loses the signal. At that point a maintained interval and a maintained signal cease to coincide.

Table 3 compares three configurations across the seven criteria that structured the coding. The comparison is not monotone. The transition from the first configuration to the second improves most criteria. The transition from the second to the third reverses that improvement while the volume of communication continues to grow, which is why the loss stays invisible in operational reporting.

Table 3: Communication without artificial intelligence, with human controlled artificial intelligence, and under full automation

Criterion	Without the technology	With human control	Under full automation	Effect on rhythm
Transparency	Limited by capacity to report	Configuration of use is stated	Disclosure becomes formal and unread	Amplitude falls
Authenticity of communication	High and unstable	Preserved through revision	Voice becomes generic	Amplitude falls
Personalization	Coarse and manual	Adapted to behavior with consent	Perceived as surveillance	Amplitude falls
Speed and timeliness	Dependent on staff	Stable and scheduled	Stable and scheduled	Frequency rises
Human control	Complete	Located at defined review points	Absent	Predictability rises, trust falls
Risk of error and loss of trust	Local and correctable	Bounded by review	Systemic and repeated at scale	Amplitude falls sharply
Effect on trust and loyalty	Slow accumulation	Accelerated accumulation	Contact grows without trust	Divergence of parameters

Destruction of trust follows an order that can be stated explicitly. Undisclosed use revealed by a third party ranks first. Personalization crossing into perceived surveillance ranks second, and its mechanism is documented, since personalization works only where trust in the collector of data already exists (Aguirre *et al.*, 2015). Targeting errors that place content in a sensitive context rank third, automated response to a crisis ranks fourth, and loss of amplitude under a preserved calendar ranks fifth as the slowest and least visible. The first two triggers damage the channel rather than the message, which is why the harm is disproportionate to the scale of the incident. Calibration applies in full here, since excessive delegation costs as much as insufficient use (Lee & See, 2004).

Recovery proceeds in five steps whose order is not interchangeable. The organization first acknowledges the incident before the facts are fully established, because attribution consolidates faster than verification. It then discloses the configuration of use, since the audience now assesses the process rather than the output. Third, it narrows or suspends the automated loop in the affected domain. Fourth, it places the content that continues to be published under human verification. Finally, it restores the ordinary cadence. This last step is the one most often abandoned once acute pressure subsides.

Task Allocation, Indicators and Propositions

Allocation follows from two properties of a communication act, the objectivity of the underlying task and its relational weight, both established in the review of task boundaries. Figure 2 presents them as the axes of a matrix. The horizontal axis runs from subjective to objective tasks, and a strip beneath it maps the four states of autonomy onto that axis, from assisted through augmented and delegated to autonomous. The vertical axis runs from low to high relational weight. A dashed line divides the matrix horizontally and marks the threshold of readable human authorship. The line shows that the threshold can be crossed in two ways: by raising the level of autonomy, or by delegating an act of higher relational weight at an unchanged level of autonomy. The four quadrants define the modes of allocation, namely full delegation, delegation with human revision, human authorship with machine support, and exclusively human authorship. Segmentation, scheduling and routine reporting fall under full delegation, while fundraising appeals, crisis statements and apologies after a failure fall under exclusively human authorship. The matrix is the most transferable element of the paper, because it converts a general principle into a decision a team can make about a specific task on a specific day (Fig. 2).

Monitoring requires indicators that register the




Figure 2: Allocation of marketing tasks between people and artificial intelligence

behavioral consequence of trust without asking about it. Three carry most of the diagnostic load: variance of the interval between publications, share of communication acts with readable human authorship, and median latency of a human reply to criticism. Two requirements govern their use. Measurement must be longitudinal, because the central variable is variance in an interval and a point estimate cannot express it. Measurement must also separate engagement from trust, since mechanics that

raise engagement can leave trust unchanged (Demydova, 2025b). The model yields six propositions, each stated with the condition that would falsify it (Table 4).

Numerical Illustration of the Frequency and Amplitude Trade Off

Propositions P1 and P4 state relationships that pull in opposite directions, so a demonstration that both can hold simultaneously is required before any empirical test is designed. Evaluating Equations 1 to 5 across the full range of α yields a curve that is not monotone. The index rises from one hundred at full human authorship to a maximum of about one hundred forty two, reached where roughly forty four percent of communication acts pass without a human review point. Beyond that point the index falls. At full autonomy it stands at about twenty seven, below the value obtained without any automation at all. Figure 3 plots the trust index as a solid curve against the share of unreviewed acts, together with its frequency component as a dotted line and its amplitude component as a dashed line, both rescaled so that they equal one hundred at zero automation. A vertical marker locates the interior maximum. The frequency component keeps rising across the entire range while the trust index collapses, which is exactly the divergence of rhythm parameters described above. The figure is essential because it makes visible that the gain in

Table 4: Propositions and falsification routes

No.	Proposition	Condition that falsifies it
P1	Regularity of communication affects trust more than its volume, and the relationship holds under machine mediation	Trust responds to volume irrespective of interval variance
P2	The effect of disclosure is mediated by perceived human control rather than by the fact of disclosure	Disclosure lowers trust equally across configurations of control
P3	Delegation of relational acts lowers trust even when the frequency of contact is unchanged	Trust is unaffected by allocation of relational acts once frequency is held constant
P4	A gain in frequency raises trust only up to the point at which the loss of amplitude exceeds it	Trust rises monotonically with frequency at every level of amplitude
P5	Undisclosed use revealed by a third party damages trust more than prior disclosure of the same fact	The two conditions produce declines of comparable depth
P6	The threshold of admissible automation is lower in nonprofit organizations than in commercial ones	The threshold is invariant across sectors under equivalent conditions




Figure 3: Simulated behavior of the trust index as the share of unreviewed communication acts increases

frequency and the loss of trust occur at the same time (Fig. 3).

Three consequences follow. The interior maximum shows that P1 and P4 are compatible: regularity raises trust, and additional regularity purchased through loss of substance lowers it. The position of the maximum is a property of the assumed parameters and not a recommendation, so no organization should read forty four percent as a target. The shape of the curve, however, is robust to the parameter values, because any specification in which amplitude declines faster than frequency grows produces an interior maximum. That shape is the testable prediction. A study that measured trust across organizations with differing shares of unreviewed communication and found a monotone relationship would falsify P4 and, with it, this part of the model.

Theoretical Position and Practical Implications

Three accounts compete for the explanation of consumer response to machine authorship. Intention based models explain the reaction of a consumer to an advertising stimulus, and evidence from online advertising shows that trust may even fail to moderate that reaction (Aktar & Gacho Segumpan, 2025). Research on aversion and appreciation explains the reaction of a person to a machine recommendation and resolves its internal contradiction through the type of task (Dietvorst *et al.*, 2015; Logg *et al.*, 2019). Research on disclosure explains the reaction of an audience to a label (Haupt *et al.*, 2025). All three explain a single act. None explains the trajectory that accumulates across a year, and the trajectory is what a marketing department actually manages. The contribution of the present model lies in that shift of unit. The position is compatible with the experiential account of consumer interaction with such systems (Puntoni *et al.*, 2021) and with chain models holding that positioning acquires force only through consistent reproduction (Balaian, 2026c). One objection deserves an answer. Amplitude may appear unmeasurable, yet it is observable indirectly through the share of communication acts carrying verifiable external evidence. Difficulty of measurement is not an argument against the existence of a variable. Four practical implications follow. Organizations with weak trust usually increase the volume of what they publish. The model instead recommends stabilizing the interval first and placing automation where regularity breaks down for resource reasons, a task for which lifecycle messaging practice supplies the vocabulary (Andreasyan, 2026e). Disclosure should report the configuration of human control rather than the bare fact of machine use, which is a design task of the same class as the design of trust signals inside a funnel (Balaian, 2026d). Relational acts remain with people even where this reduces throughput, and an omnichannel architecture makes that feasible by letting the routine layer run automatically (Avagimyan, 2026c).

Engagement metrics and trust indicators are reported separately, and reporting standards that make attribution logic explicit provide the format for that separation (Avagimyan, 2026b; Balaian, 2026a).

Boundary Conditions in Nonprofit and Service Organizations

The model applies to commercial and nonprofit organizations alike, but the threshold within it moves. The three structural features of the nonprofit environment identified in the literature review explain the movement. A contributor who cannot verify the result relies on the credibility of the account rather than on a product received. The separation of payer and beneficiary weakens every analogy borrowed from commercial personalization, since the recipient of the message is not the source of revenue. The practical consequence is stated in P6. A configuration that is safe in a commercial funnel becomes risky in a fundraising appeal, and the same message that reads as convenient automation from a retailer reads as indifference from a charity. The evidence supporting this claim is thin (Cheng & Wang, 2025), which is why it appears as a boundary condition and a proposition rather than as a result.

Professional service organizations occupy an intermediate position between the two poles and fall under the same logic. The result of an advisory engagement is delayed and partly unobservable, so a client extends credit against future conduct in a way that resembles the position of a donor. Key performance indicator systems built around trust, outcomes and renewal describe exactly this dependence (Avagimyan, 2026e), and trust signals inside a consulting funnel are engineered rather than incidental (Balaian, 2026d). The threshold of admissible automation in such organizations sits below the commercial retail level and above the nonprofit level. All three sectors therefore lie on a single scale rather than in separate categories.

Limitations and Directions for Further Research

The model is untested. Every proposition in Table 4 states a relationship that the reviewed literature makes plausible, and none has been estimated on data collected for this purpose. The simulation illustrates internal consistency and is not evidence, since the authors chose its parameters rather than estimating them from observation. The limitations of the method apply to the results in full, and the thin evidence on nonprofit communication makes the boundary conditions for that sector the least secure part of the paper.

Four directions follow. Panel measurement of publication timing matched to periodic assessment of brand trust would test P1 and P4 and requires only archival collection of dates paired with an existing survey series. An experiment varying the configuration of disclosure while holding the fact of machine use constant would test P2, and testing protocols used

in applied digital marketing research offer a template (Demidova, 2025a). A field study reallocating relational acts while holding contact frequency constant would test P3. Comparative work across commercial and nonprofit organizations would test P6. It carries the highest value for practice, since the sector that most depends on trust currently operates with the least evidence about it, and the diagnostic instruments already used for lead qualification and growth bottlenecks can structure such comparisons (Andreasyan, 2026a; Balaian, 2026b).

CONCLUSION

This paper reduced the documented antecedents of trust in machine mediated marketing to four operational domains, namely transparency, authenticity, timeliness and human accountability, and attached to their enactment three parameters of rhythm, namely frequency, amplitude and predictability. Trust is sustained by the periodic reproduction of these domains rather than by their presence, and artificial intelligence acts on the parameters in opposite directions at the same time.

Three further results follow from that architecture. Automation raises frequency and predictability because it removes the dependence of the communication calendar on staff availability, and it lowers amplitude because volume without organizational substance carries no verifiable information. The decisive variable in the disclosure debate is not the label attached to a message but the readability of human authorship behind it, which locates the threshold between delegated and autonomous operation. Destruction of trust proceeds hierarchically, and the triggers that damage the evidentiary channel cause harm disproportionate to the scale of the incident. The study contributes a matrix that allocates marketing tasks by the objectivity of the task and the relational weight of the act, a set of indicators for continuous monitoring of rhythm, and six propositions stated with explicit falsification routes. Its principal limitation is that the framework is theoretical and rests on evidence concentrated in commercial settings, so its value depends on longitudinal and experimental testing, with comparative work across sectors as the most urgent recommendation. The practical claim it supports is narrower than the usual prescription and more demanding. Trust is not accumulated through the volume of communication and is not defended through speed. It is held only while a person remains visible inside a rhythm that the audience can predict.

REFERENCES

- 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>
- Aktar, L., & Gacho Segumpan, R. (2025). Online advertisement and online purchase intention among Bangladeshi domestic tourists: Trust as a moderator. *American Journal of Economics and Business Innovation*, 4(3), 61-68. <https://doi.org/10.54536/ajebi.v4i3.5720>
- Andreasyan, A. (2026a). Algorithm for lead qualification and proposal personalization: A scoring model for marketing consulting sales funnels. *Current Research Bulletin*, 3(4), 122-129. <https://doi.org/10.55677/CRB/I4-04-CRB2026>
- Andreasyan, A. (2026b). Methodology for marketing audit and growth prioritization in SMEs: KAPRI method, a KPI-linked diagnostic framework. *Journal of Business and Strategic Management*, 11(4), 44-59. <https://doi.org/10.47941/jbsm.3801>
- Andreasyan, A. (2026c). Strategy VPR-Align model for positioning and differentiation in marketing consulting: An empirical framework linking value proposition to revenue model. *Journal of Public Relations*, 5(1), 1-14. <https://doi.org/10.47941/jpr.3665>
- Andreasyan, A. (2026d). Toward sustainable ROI governance in marketing: Integrating multi-touch attribution and incrementality testing for evidence-based budget decisions. *Economics, Management and Sustainability*, 11(1), 218-227. <https://doi.org/10.14254/jems.2026.11-1.12>
- Andreasyan, A. (2026e). Ways to increase customer retention: Practical tactics for lifecycle messaging, loyalty mechanics, and churn prevention. *Current Research Bulletin*, 3(6), 191-199. <https://doi.org/10.55677/CRB/I6-02-CRB2026>
- Avagimyan, A. (2026a). Methodology for diagnosing client growth bottlenecks in marketing consulting: An empirical mixed-methods toolkit. *International Journal of Recent Trends in Multidisciplinary Research*, 6(4), 126-136. <https://doi.org/10.59256/ijrtmr.20260604014>
- Avagimyan, A. (2026b). Methods and approaches to prove ROI of marketing consulting engagements: An empirical attribution protocol for advisory services. *Current Research Bulletin*, 3(6), 206-211. <https://doi.org/10.55677/CRB/I6-04-CRB2026>
- Avagimyan, A. (2026c). Strategy for omnichannel brand communication: OBRIDGE model, evidence-based segmentation, personalization, and retention design. *Journal of Entrepreneurship and Project Management*, 11(2), 54-66. <https://doi.org/10.47941/jepm.3815>
- Avagimyan, A. (2026d). Sustainable human capital allocation in professional services: An integrated predictive scoring architecture for capacity-constrained lead governance. *Economics, Management and Sustainability*, 11(2), 6-16. <https://doi.org/10.14254/jems.2026.11-2.1>
- Avagimyan, A. (2026e). Ways to measure and improve client retention in marketing consulting: Tor-Bridge model and an empirical KPI system for trust, outcomes, and renewal. *Current Research Bulletin*, 3(4), 130-138. <https://doi.org/10.55677/CRB/I4-05-CRB2026>
- Baek, T. H., Kim, J., & Kim, J. H. (2024). Effect of disclosing AI-generated content on prosocial

- advertising evaluation. *International Journal of Advertising*, 45(1), 171-192. <https://doi.org/10.1080/02650487.2024.2401319>
- Balaian, A. (2026a). Approaches to measuring marketing consulting impact trace-map method: KPIs, attribution logic, and client reporting standards. *Current Research Bulletin*, 3(4), 107-114. <https://doi.org/10.55677/CRB/14-02-CRB2026>
- Balaian, A. (2026b). Methodology for diagnosing client growth bottlenecks in marketing consulting: A data-informed framework. *Journal of Business and Strategic Management*, 11(3), 82-95. <https://doi.org/10.47941/jbsm.3683>
- Balaian, A. (2026c). Strategy for positioning and differentiation in marketing consulting V-FLOW chain model: From value proposition to revenue model. *Journal of Entrepreneurship and Project Management*, 11(2), 1-15. <https://doi.org/10.47941/jepm.3681>
- Balaian, A. (2026d). Ways to increase consulting engagement conversion in marketing: Trust signals, offers, and funnel design. *Human Resource and Leadership Journal*, 11(2), 60-73. <https://doi.org/10.47941/hrlj.3682>
- Balayan, A. (2026). An algorithmic framework for sustainable marketing audit and prioritization in consulting engagements: The PRISM-Bridge model. *Economics, Management and Sustainability*, 11(1), 136-145. <https://doi.org/10.14254/jems.2026.11-1.7>
- Castelo, N., Bos, M. W., & Lehmann, D. R. (2019). Task-dependent algorithm aversion. *Journal of Marketing Research*, 56(5), 809-825. <https://doi.org/10.1177/0022243719851788>
- Cheng, Y., & Wang, Y. (2025). Leveraging artificial intelligence-powered chatbots for nonprofit organizations: Examining the antecedents and outcomes of chatbot trust and social media engagement. *Journal of Philanthropy and Marketing*, 30, Article e70013. <https://doi.org/10.1002/nvsm.70013>
- Demidova, A. (2025a). Artificial intelligence in digital marketing: Methodology, advantages, and constraints. *Economic Bulletin of Dnipro University of Technology*, 2025(3), 34-44. <https://doi.org/10.33271/ebdut/91.034>
- Demidova, A. (2025b). Integration of automated tools in the lead generation process: Chatbots, autofunnels and application collection systems. *Journal of Bio-based Marketing*, 3. <https://doi.org/10.5281/zenodo.15909977>
- Demidova, A. (2026). Methodology for attracting clients through social networks in the services niche: From content strategy to application. *SBS Journal of Applied Business Research*, 14(1). <https://doi.org/10.70301/JOUR/SBS-JABR/2026/14/1/1>
- Demydova, A. (2025a). Lead generation strategy based on Meta platforms: User behavior, warming up, and conversion. *International Journal of Education, Business and Economics Research*, 5(5), 1-13. <https://doi.org/10.59822/IJEBER.2025.5501>
- Demydova, A. (2025b). Level-up loyalty: Gamified engagement and client lifetime value in online service markets. *International Journal of Accounting Finance and Social Science Research*, 3(6), 185-198. <https://doi.org/10.63452/IJAFSSR.2025.3613>
- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. *Journal of Experimental Psychology: General*, 144(1), 114-126. <https://doi.org/10.1037/xge0000033>
- Dietvorst, B. J., Simmons, J. P., & Massey, C. (2018). Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science*, 64(3), 1155-1170. <https://doi.org/10.1287/mnsc.2016.2643>
- Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*, 14(2), 627-660. <https://doi.org/10.5465/annals.2018.0057>
- Haupt, M., Freidank, J., & Haas, A. (2025). Consumer responses to human-AI collaboration at organizational frontlines: Strategies to escape algorithm aversion in content creation. *Review of Managerial Science*, 19(2), 377-413. <https://doi.org/10.1007/s11846-024-00748-y>
- Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1), 50-80. https://doi.org/10.1518/hfes.46.1.50_30392
- Leung, E., Paolacci, G., & Puntoni, S. (2018). Man versus machine: Resisting automation in identity-based consumer behavior. *Journal of Marketing Research*, 55(6), 818-831. <https://doi.org/10.1177/0022243718818423>
- 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. <https://doi.org/10.1016/j.obhdp.2018.12.005>
- Longoni, C., Bonezzi, A., & Morewedge, C. K. (2019). Resistance to medical artificial intelligence. *Journal of Consumer Research*, 46(3), 407-434. <https://doi.org/10.1093/jcr/ucz013>
- Luo, X., Tong, S., Fang, Z., & Qu, Z. (2019). Machines versus humans: The impact of AI chatbot disclosure on customer purchases. *Marketing Science*, 38(6), 937-947. <https://doi.org/10.1287/mksc.2019.1192>
- Mayer, R. C., Davis, J. H., & Schoorman, F. D. (1995). An integrative model of organizational trust. *Academy of Management Review*, 20(3), 709-734. <https://doi.org/10.2307/258792>
- Ozikor, A. F., Mani, E. G. O., & Deborah, B. A. (2026). Artificial intelligence in PR and advertising campaigns: Ethical and strategic implications. *Journal of Media, Journalism & Mass Communication*, 2(1), 74-79. <https://doi.org/10.54536/jmjmc.v2i1.6484>
- 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>