


**Tourism Policy**

ISSN: 3149-8167

available at http: www.econjournals.com

**Tourism Policy, 2026, 2(4), 9-17.**

# Regulating Generative Artificial Intelligence Travel Agents to Protect Tourist Consumers from Algorithmic Misinformation

**Muddassar Sarfraz***

College of Business Administration, Prince Mohammad Bin Fahd University, Al-Khobar, Saudi Arabia.
*Email: msarfraz@pmu.edu.sa

**Received:** 30 July 2026 **Accepted:** 04 November 2026 **DOI:** https://doi.org/10.32479/tp.25351

**ABSTRACT**

**G**enerative AI travel agents can combine useful advice with plausible inaccuracies, raising questions about the protection afforded by disclosure and consumer choice. This study integrates a secondary analysis of a public survey released in 2023 with a focused assessment of European Union and United Kingdom regulation. The dataset contains 900 consumer records; 886 complete cases enter the adjusted models. Among respondents whose mean perceived hallucination-potential score was at least four on a five-point scale, 58 of 224 endorsed future AI travel-planning use (25.9%; 95% confidence interval 20.6-32.0%). The standardized risk–intention association changed from 0.181 without controls to −0.049 after adjustment for experience and acceptance-related perceptions. These observational estimates establish neither misinformation detection nor the effectiveness of warnings. The legal assessment distinguishes AI-origin transparency from the accuracy and commercial presentation of travel advice. The paper proposes proportionate verification of consequential claims, confirmation before financial commitments, accessible correction, and an identifiable trader responsible for complaint handling. Tourism authorities should evaluate factual reliability and recoverability alongside disclosure compliance, while testing these safeguards prospectively.

**Keywords:** Generative Artificial Intelligence, Travel Agents, Consumer Protection, Algorithmic Misinformation, Tourism Governance, Regulatory Policy
**JEL Classifications:** L83, M38, Z3

## 1. INTRODUCTION

A travel recommendation becomes consequential when a tourist uses it to buy a ticket, select accommodation, or decide whether a journey is feasible. A conversational system may present an attraction as open, describe an unverified cancellation entitlement, or combine a plausible route with an impossible connection. The fluent presentation of such claims can obscure the distinction between an attractive suggestion and information on which a consumer can safely act. Generative AI therefore raises a practical tourism-policy question: which duties should attach to the commercial provision of travel advice when persuasive language can exceed the reliability of its factual basis?

Recent tourism scholarship identifies applications of large language models across planning, customer service, and business operations, together with risks arising from inaccurate outputs and uncertain information provenance (Carvalho and Ivanov, 2024). Experimental research now documents adverse responses to hallucinated itineraries and service failures, including changes in perceived accuracy, trustworthiness, loyalty, and word-of-mouth intentions (Belanche et al., 2026; Rejón-Guardia et al., 2026). These findings make factual reliability relevant to tourism governance as well as service management. A failed recommendation may affect a visitor's expenditure or ability to access a service even when the underlying booking system functions correctly.

This Journal is licensed under a Creative Commons Attribution 4.0 International License

Consumer awareness is one possible safeguard, but awareness must be interpreted carefully. Knowing that AI can be wrong does not establish whether a person can recognize a particular error, verify a claim in another language, or recover a loss after relying on it. Intention to use a system also differs from uncritical reliance. A consumer may knowingly use AI to generate ideas while checking prices elsewhere. Regulation needs evidence that respects these distinctions rather than equating adoption, trust, and vulnerability.

This paper examines the coexistence of perceived misinformation potential and intended AI use through a transparent reanalysis of an openly available tourism-consumer survey (Hansen et al., 2023). It then assesses how existing transparency and unfair-trading rules address the relevant commercial conduct. The jurisdictions are the European Union (EU) and the United Kingdom (UK), selected because they provide contrasting, documented routes through AI-specific and general consumer regulation. The analysis concerns the legal position as of October 25, 2026. It does not treat these jurisdictions as an exhaustive account of international tourism law.

Three research questions guide the study. How frequently do stronger beliefs about AI hallucination potential coexist with intended use for travel decisions? How does their association change after accounting for familiarity, prior use, and perceived benefits? Which regulatory obligations and additional tourism-specific safeguards address the resulting information problem? The empirical analysis answers the first two questions. The third requires legal interpretation and normative judgment; it is not estimated from the survey.

The contribution is to connect consumer-level evidence with the design of verifiable commercial duties. Earlier work using the same survey examined technology acceptance and itinerary choices (Christensen et al., 2025). The present study uses only the variables actually deposited in the public archive, reports absolute co-occurrence frequencies, examines sensitivity to the content of the risk measure, and develops a separate regulatory assessment. Its evidence supports a limited conclusion: high perceived error potential can coexist with intended use, so policy should address what happens at the point of consequential reliance as well as whether consumers recognize AI involvement.

## 2. LITERATURE REVIEW

### 2.1. Algorithmic Misinformation in Tourism Services

In this paper, algorithmic misinformation means materially false, misleading, or unsupported travel assertions generated or relayed through an automated system. The definition concerns the information presented to the consumer and does not presume an intention to deceive. Hallucination refers more narrowly to generated content that lacks an adequate factual basis or conflicts with the relevant source or context. Research on natural-language generation distinguishes several forms of hallucination and shows that evaluation depends on the task and its relationship to source information (Ji et al., 2023).

The wider definition matters in tourism because an incorrect answer can arise through different mechanisms. A fabricated attraction, an outdated opening time, an omitted mandatory charge, and a false claim of refund eligibility can produce similar consumer difficulties while requiring different remedies. Retrieval from a current source may reduce one problem without resolving another. A correct source may be summarized inaccurately, and a factually accurate price may still be presented misleadingly if unavoidable charges are hidden. An adequate regulatory assessment must therefore examine both information quality and commercial presentation.

The term travel agent is used functionally here for a consumer-facing system that helps formulate travel choices. It covers conversational planning and recommendations, assistance with booking, and systems authorized to execute transactions. These functions imply different levels of potential reliance. The survey examined in this study concerns planning intentions in 2023. It provides no direct evidence about autonomous purchasing agents. Recommendations concerning transaction execution are consequently reasoned extensions based on the consequences of delegated action.

### 2.2. Consumer Acceptance and Calibrated Reliance

Acceptance research explains why people find AI-assisted planning attractive. Topsakal (2025) examines familiarity, ease of use, usefulness, and trust in travel-planning acceptance. These constructs describe perceived value and accessibility; they do not establish objective output quality. Responsible-AI research in tourism similarly connects ethical design with the conditions under which consumers place trust in a system (Singu et al., 2026). A policy assessment should preserve the benefits of accessible planning assistance while asking whether reliance is appropriate to the reliability of the particular task.

The empirical literature cautions against assuming that an error has a uniform behavioral effect. Rejón-Guardia et al. (2026) find that hallucinations reduce perceived accuracy and can affect itinerary acceptance through a sequence involving usefulness and trustworthiness. Their experiments also identify the importance of outcomes and error salience. Belanche et al. (2026) distinguish hallucinations from ordinary service mistakes and examine their consequences for customer relationships. Together, these studies suggest that general beliefs about risk, the recognition of a specific failure, and subsequent responses belong to different stages of the consumer process.

Disclosure research also addresses several distinct objects. An AI label identifies a source or mode of production. A warning communicates a limitation. A factual citation permits checking. An explanation of a booking condition specifies a commercial entitlement. Their effects should not be assumed interchangeable. Balaban and Phua (2026), studying travel-destination advertising, find that disclosure can help consumers recognize AI-generated content, with responses depending on the source context. Huang et al. (2026) examine creator identity disclosure and perceptions of responsibility in tourism management. Neither result makes an AI-origin label a guarantee that a recommended fare, opening time, or connection is correct.

Sun et al. (2026) examine travel persuasion in relation to recommender characteristics, message strategies, and hallucination reminders. This task-sensitive approach is relevant to consumer protection: A notice that works for an inspirational destination message may work differently when the consumer is choosing a nonrefundable product. The present study does not test any notice. It uses observed risk beliefs and intentions to examine a premise that sometimes enters policy discussion, namely that consumers who appreciate the potential for misinformation will necessarily avoid AI-assisted travel decisions.

### 2.3. Responsibility and Consumer Protection

Responsible AI extends beyond increasing willingness to adopt a service. Shen et al. (2026) conceptualize responsible generative AI in tourism through authenticity, security, caring, ethicality, and fairness. These dimensions help locate information reliability within a wider account of the tourist experience. For regulatory purposes, however, principles must be connected to conduct that a firm can document and an authority can assess. Clarke (2025) emphasizes the relationship between technological characteristics, patterns of use, and responsibility for adverse consequences.

Consumer law provides one route for making that connection. Duivenvoorde (2025) examines synthetic advertising, personalized persuasion, and consumer-facing generative chatbots, identifying both protections and limitations in the EU framework. This scholarship challenges the assumption that introducing a generative model necessarily creates a complete legal vacuum. At the same time, established rules do not automatically provide a sector-specific verification workflow, a complete allocation of responsibility across suppliers, or an easily usable remedy for a tourist abroad.

The analytical distinction adopted here is between origin transparency, substantive commercial accuracy, and practical recoverability. Each addresses a different question: whether the tourist knows AI is involved, whether a material claim can reasonably support a decision, and whether an error can be corrected or challenged. The proposed safeguards are evaluated against these functions. Their selection is informed by consumer evidence and legal coverage, while their effectiveness remains an empirical question for future field studies.

## 3. MATERIALS AND METHODS

### 3.1. Research Design and Data Source

The study combines secondary quantitative analysis with a focused doctrinal assessment. The quantitative component uses the Technology Acceptance Model (TAM) and AI Hallucination Survey Dataset, deposited at Utah State University on September 05, 2023 under a Creative Commons Attribution 4.0 license (Hansen et al., 2023). Data were collected through Qualtrics in May-June 2023. The downloaded comma-separated file contains 900 records, 43 rating items, and a response identifier.

The archive supplies no demographic fields, country indicators, sampling weights, itinerary selections, or open-ended responses. The analysis therefore cannot reproduce the choice component of the associated publication or estimate population-representative rates. The 900 records are treated as the available analytical sample. No new participants were recruited. Response identifiers were removed from the working analysis, and no attempt was made to identify individuals.

### 3.2. Measurement and Data Preparation

Table 1 describes the deposited item blocks. Higher scores indicate stronger intention, familiarity, perceived benefits, comfort with AI performing travel tasks, or perceived hallucination potential, according to the block. Prior use measures frequency. Each composite is the unweighted mean of its deposited items, calculated only when every item in that block is observed:

$$S_{ic} = \frac{1}{K_c} \sum_{j=1}^{K_c} x_{icj} \quad (1)$$

In Equation (1), S is the score for respondent i and construct c, x is an item response, and K is the number of deposited items in the block. This rule avoids assigning a score to an incompletely answered block in the main analysis. Composite construction does not involve estimated factor weights or reconstruction of unobserved survey items.

All nonmissing responses in the deposited file are integers from one to five. The repository README describes seven-point scales, whereas the associated article specifies five-point measures (Christensen et al., 2025). The analysis retains the observed coding without rescaling. The archived usefulness and comfort blocks contain four and seven columns, respectively, fewer than the entries listed for these constructs in the published measurement table. These controls are consequently interpreted as the deposited blocks rather than reconstructions of the longer published instruments. Their complete item-to-wording correspondence cannot be established from column labels alone.

The five-item hallucination-potential block covers fabricated information, imaginary scenarios, cultural bias, flawed predictions, and manipulation or hacking. It measures perceived likelihood across related risks, rather than tested factual knowledge. A narrower sensitivity measure averages the first two items, which address false information and fabrication. The first intention item

**Table 1: Descriptive statistics and internal consistency of the archived scales**

| Construct | Items | n | Mean | SD | α |
|---|---|---|---|---|---|
| Intention to use AI | 7 | 900 | 2.496 | 1.130 | 0.956 |
| Hallucination potential | 5 | 900 | 3.160 | 0.992 | 0.928 |
| AI familiarity | 5 | 900 | 3.008 | 1.199 | 0.935 |
| Prior general use | 6 | 900 | 2.130 | 1.167 | 0.937 |
| Perceived usefulness | 4 | 900 | 2.987 | 1.131 | 0.953 |
| Perceived ease of use | 5 | 900 | 2.868 | 1.054 | 0.943 |
| Attitude toward use | 4 | 900 | 3.023 | 1.079 | 0.938 |
| Comfort delegating tasks | 7 | 886 | 3.161 | 1.020 | 0.949 |

Source: Calculations from Hansen et al. (2023). Scores are unweighted item means on the observed 1-5 coding. The comfort score requires seven valid answers. α denotes Cronbach's alpha, calculated on complete item blocks; SD is the sample standard deviation. The risk block measures perceived likelihood, not factual accuracy or error-detection ability

concerns future travel-planning use; its fourth and fifth response categories define direct endorsement. A separate item concerns considering AI for checking travel prices. The full seven-item intention composite includes a wider range of tourism decisions and recommendation intentions.

Whitespace-only cells were converted to missing values. There are 30 such cells, confined to the comfort block and distributed across 14 respondents. All other blocks are complete. Descriptive analyses of the focal variables use 900 records. The main nested regressions use the same 886 complete cases, ensuring that changes across specifications do not reflect changes in the estimation sample.

### 3.3. Statistical Analysis

The descriptive analysis reports the joint occurrence of intended use and perceived hallucination potential. Risk-score thresholds of three and four identify, respectively, averages at least at the moderately likely and very likely response levels. These thresholds describe response levels rather than empirically validated consumer segments. They are nested thresholds. The primary co-occurrence outcome is direct intention endorsement. Price-checking endorsement and an intention-composite threshold of four provide complementary summaries. Conditional proportions use the relevant risk group's denominator, with 95% Wilson intervals. The intervals describe sampling uncertainty under independent observations and do not address selection into the original survey.

Ordinary least squares estimates the association between perceived hallucination potential and the intention composite. All scores are standardized using the mean and sample standard deviation of the 886 complete cases. The full specification is:

$$B_i = \alpha + \beta H_i + \gamma F_i + \delta U_i + X_i'\theta + \varepsilon_i \quad (2)$$

In Equation (2), B denotes standardized intention, H hallucination potential, F familiarity, U prior use, and X the standardized usefulness, ease, attitude, and delegation-comfort scores. The coefficient on H describes a partial association. Model 1 includes H alone; Model 2 adds F and U; Model 3 adds X. The intercept is retained, although it is zero to numerical precision after standardization. The models organize descriptive associations; they do not identify the causal effect of awareness or a mediated effect through acceptance perceptions.

Inference uses HC3 heteroskedasticity-consistent standard errors and Student t critical values with residual degrees of freedom. A respondent bootstrap with 4,999 resamples and seed 20261025 re-estimates Model 3, repeating standardization in each resample. Reported bootstrap bounds are percentile intervals. Model fit and variance-inflation factors are descriptive diagnostics. Since item responses are ordinal, a sensitivity model replaces every score with its average rank before standardization.

Further checks use the narrower two-item risk score, individual risk items, the direct intention item as the outcome, and alternative treatment of incomplete comfort responses. One model omits comfort and retains all 900 records; another averages comfort when at least six of seven answers are available. An exploratory interaction between standardized risk and familiarity is also examined. Sensitivity analyses are not independent confirmatory tests, and their intervals are pointwise without multiplicity adjustment. All calculations use the deposited numerical observations rather than values copied from the original article.

### 3.4. Legal Assessment

The legal component examines provisions directly relevant to consumer-facing travel information in two jurisdictions. The source set comprises Article 50 of the EU Artificial Intelligence Act, its current presentation by the European Commission and associated transparency guidance, the UK Digital Markets, Competition and Consumers Act 2024, and the Competition and Markets Authority's unfair-commercial-practices guidance (Competition and Markets Authority [CMA], 2025; Digital Markets, Competition and Consumers Act 2024; European Commission, 2026a, 2026b). The assessment was closed on October 25, 2026. It is a purposive comparison, not a systematic review of every potentially applicable legal instrument.

Each source is examined for the actor addressed, the triggering activity, the relevant duty, and the aspect of consumer protection it covers. Statutory requirements are distinguished from interpretive guidance and the paper's policy proposals. The EU source establishes transparency obligations; the UK sources address commercial conduct. Cross-jurisdictional differences are retained rather than collapsed into a single universal rule. The final recommendations are reasoned proposals judged by their fit with the identified information problem, expected administrative burden, and capacity to produce reviewable evidence.

## 4. RESULTS

### 4.1. Measurement and Co-occurrence

Mean hallucination potential is 3.160 (SD 0.992), compared with 2.496 (SD 1.130) for intention (Table 1). Internal-consistency coefficients range from 0.928 to 0.956 across the deposited blocks. This consistency supports averaging their items but does not demonstrate objective knowledge, independence of constructs, or absence of common-method bias. Figure 1 shows substantial variation in both focal scores and the joint occurrence of high scores.

Of the full sample, 176 respondents endorse direct intention to use AI for future travel planning, and 243 endorse considering it for price checking. The hallucination-potential score is at least three for 612 respondents and at least four for 224. Among the latter group, 58 endorse direct intention, giving 25.9% (95% CI 20.6-32.0%). The corresponding number for price checking is 73, or 32.6% (26.8-39.0%). These are conditional percentages within the 224-person group; they represent 6.4% and 8.1% of all 900 records, respectively.

The more restrictive intention-composite threshold yields lower co-occurrence, as expected from its broader content: 37 of the 224

high-risk-score respondents have an intention mean of at least four, or 16.5% (12.2-21.9%). Table 2 gives the full set of summaries. The direct intention proportion is 16.3% in the group below three, 18.3% from three to below four, and 25.9% at four or above. This descriptive pattern does not support the proposition that stronger general risk beliefs necessarily coincide with avoidance.

The result is not confined to averaging the broad five-item risk block. Among 342 respondents selecting one of the top two likelihood categories on the first risk item, which concerns false information presented as fact, 86 endorse direct intention: 25.1% (95% CI 20.8-30.0%). These respondents may intend to verify outputs or restrict their use to low-consequence tasks. The archive contains no observations that would distinguish those possibilities from uncritical reliance.

## 4.2. Adjusted Associations and Sensitivity

Across all 900 records, the Pearson correlation between the focal composites is 0.186 and the Spearman correlation is 0.143. On the common 886-case regression sample, Model 1 produces a standardized risk coefficient of 0.181 (95% CI 0.109-0.252). It declines to 0.129 (0.064-0.193) when familiarity and prior use enter Model 2. After the acceptance-related perceptions enter Model 3, the coefficient is −0.049 (−0.093-−0.005; P = 0.028). The respondent-bootstrap interval is similar, −0.093-−0.005 (Table 3).

Standardized OLS coefficients with 95% HC3 confidence intervals in brackets; Student t critical values use residual degrees of freedom. All models use the same 886 records and include an intercept. A dash indicates that the regressor is omitted. Model 3 adds usefulness, ease, attitude, and delegation comfort. Estimates are conditional associations and do not identify causal or mediation effects.

The adjusted risk coefficient is small in standardized units. By comparison, Model 3 coefficients are 0.437 for perceived usefulness, 0.197 for attitude, and 0.159 for prior use. The model's

**Figure 1:** Distributions and the coexistence of perceived risk and intended use (n=900). Panels (a) and (b) show distributions of composite means; dashed lines mark sample means. In Panel (c), point area is proportional to the number of respondents at each score pair; dotted lines mark scores of four. Panel (d) reports direct-intention endorsement within three disjoint risk-score groups, with 95% Wilson intervals and numerator/denominator labels. Direct intention uses the first intention item, rather than the composite plotted in Panels (b) and (c)

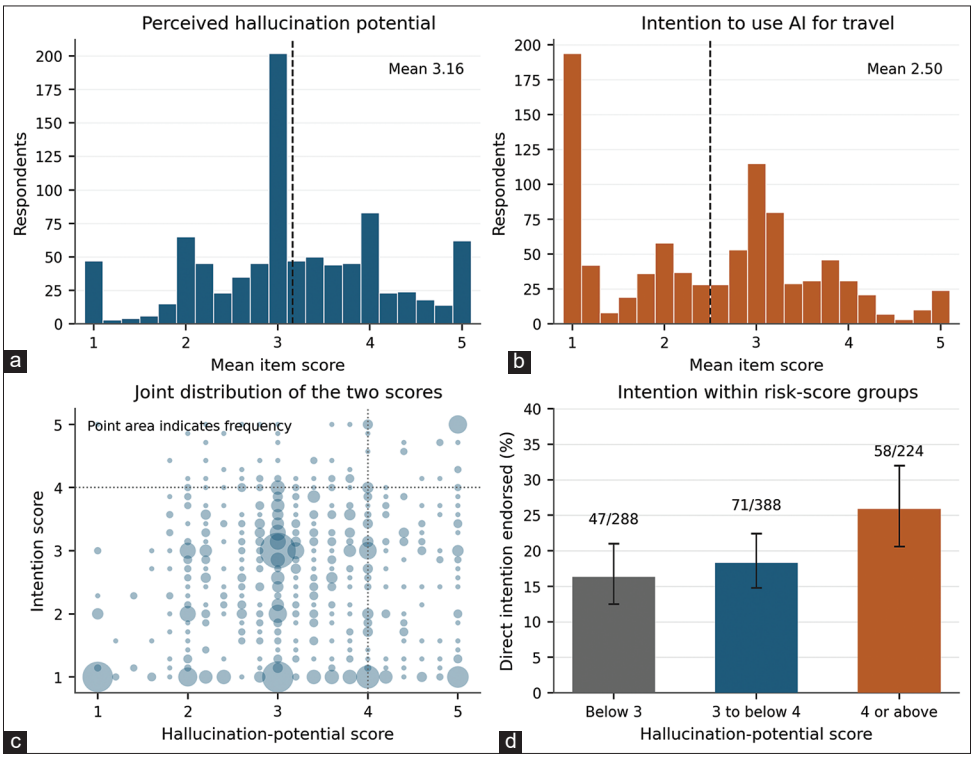



**Table 2: Intended use among respondents reporting higher hallucination potential**

| Risk-score threshold | Risk-group n | Intention criterion | Count | Within-group % | 95% CI (%) |
|---|---|---|---|---|---|
| ≥3 | 612 | Direct planning intention ≥4 | 129 | 21.1 | 18.0-24.5 |
| ≥3 | 612 | Price-checking intention ≥4 | 177 | 28.9 | 25.5-32.6 |
| ≥3 | 612 | Intention composite ≥4 | 70 | 11.4 | 9.2-14.2 |
| ≥4 | 224 | Direct planning intention ≥4 | 58 | 25.9 | 20.6-32.0 |
| ≥4 | 224 | Price-checking intention ≥4 | 73 | 32.6 | 26.8-39.0 |
| ≥4 | 224 | Intention composite ≥4 | 37 | 16.5 | 12.2-21.9 |

Source: calculations from Hansen et al. (2023), n=900. Risk thresholds refer to the five-item mean and are nested. Direct planning and price-checking criteria use their respective single items; categories 4 and 5 indicate endorsement. Confidence intervals use the Wilson method. These descriptive groups do not identify observed reliance, misinformation exposure, or consumer loss

**Table 3: Associations with the intention composite on the common complete-case sample**

| Regressor | Model 1 risk only | Model 2 plus experience | Model 3 full adjustment |
|---|---|---|---|
| Hallucination potential | 0.181 (0.109, 0.252) | 0.129 (0.064, 0.193) | −0.049 (−0.093, −0.005) |
| AI familiarity | — | 0.105 (0.023, 0.187) | −0.015 (−0.076, 0.046) |
| Prior general use | — | 0.299 (0.219, 0.380) | 0.159 (0.102, 0.217) |
| Perceived usefulness | — | — | 0.437 (0.358, 0.516) |
| Perceived ease of use | — | — | 0.074 (0.003, 0.144) |
| Attitude toward use | — | — | 0.197 (0.112, 0.283) |
| Comfort delegating tasks | — | — | 0.084 (0.025, 0.144) |
| Observations | 886 | 886 | 886 |
| $R^2$ | 0.033 | 0.171 | 0.572 |

**Table 4: Existing coverage and additional safeguards proposed for tourism services**

| Legal provision | Actor or activity | Current coverage | Additional safeguard proposed |
|---|---|---|---|
| EU Article 50(1) | Provider of a directly interacting AI system | AI-interaction notice, subject to specified exceptions | Identify the commercial service and explain its planning or execution role |
| EU Article 50(2) | Provider of a system generating synthetic output | Machine-readable marking and detectability, with exceptions and transition rules | Support consequential factual claims with appropriate current sources |
| UK ss. 226-227 | Trader engaged in a commercial practice | Misleading actions and omissions assessed under the statutory tests | Retain evidence supporting material claims and offer a correction route |
| UK ss. 229-230 | Relevant trader and invitation-to-purchase activity | Professional diligence and required purchase information | Verify key booking details and confirm the consumer's authorized commitment |

Sources: Digital Markets, Competition and Consumers Act 2024; CMA (2025); European Commission (2026a, 2026b). The final column contains this paper's proposals, not a statement that each measure is already legally mandatory. Exact duties depend on the service, actor, jurisdiction, and relevant legal test

$R^2$ is 0.572, and variance-inflation factors range from 1.13 to 3.01. The shift in the risk coefficient reflects the conditional comparison being made after correlated perceptions are included. It does not demonstrate that usefulness causes consumers to ignore risk, or that a warning would reduce intended use by a corresponding amount.

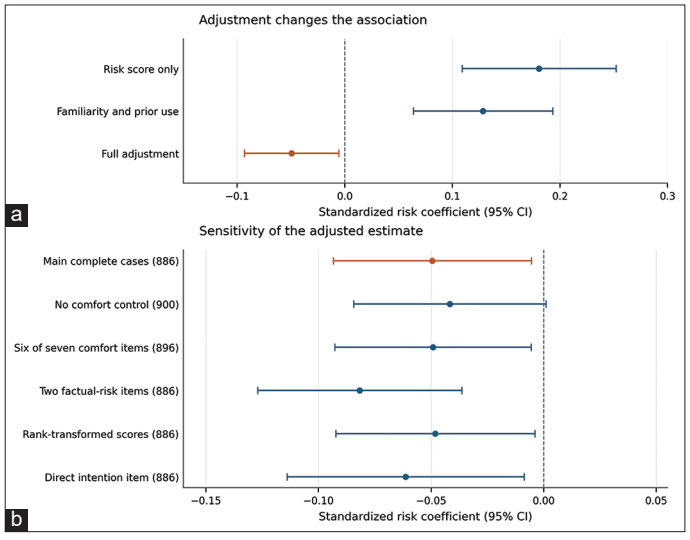
Figure 2 reports the adjustment sequence and key sensitivity checks. The narrower factual-risk measure has a coefficient of −0.082 (95% CI −0.127-−0.036). Using average ranks gives −0.048 (−0.092-−0.004), and the direct intention item gives −0.061 (−0.114-−0.009). Allowing a comfort score with six observed items retains 896 respondents and gives −0.049 (−0.093-−0.006). Omitting comfort permits all 900 records but yields −0.042 (−0.084-0.001), showing that conventional significance is not uniform across specifications.

Individual-item checks reveal additional measurement sensitivity. The risk items addressing false information and fabrication have negative adjusted coefficients of −0.064 and −0.089, respectively; both pointwise intervals exclude zero. The intervals for the items addressing bias, flawed predictions, and manipulation include zero. The exploratory risk-by-familiarity interaction is −0.001 (95% CI −0.036-0.034). There is no clear evidence in this specification that familiarity changes the conditional risk–intention relationship. These checks favor a measured interpretation centered on co-occurrence and small conditional associations rather than a uniform protective effect of perceived risk.

## 4.3. Regulatory Coverage

The EU framework addresses AI-origin transparency directly. Article 50(1) requires providers of systems intended to interact directly with natural persons to ensure that people are informed of

**Figure 2:** Adjustment and sensitivity of the risk–intention association. Points show standardized hallucination-potential coefficients with 95% pointwise HC3 intervals. Panel (a) uses the same 886 complete cases. Panel (b) shows the main model and alternative specifications; numbers in parentheses indicate sample size. The six-item comfort rule means at least six observed items. All checks retain the other full-model covariates unless specified. The dashed line marks zero. Alternative risk and outcome measures change the corresponding estimated




that interaction, subject to the provision's contextual obviousness exception. Article 50(2) separately addresses machine-readable marking and detectability of synthetic outputs. Article 50(4) contains specified disclosure duties for deepfakes and certain public-interest text (Table 4); it is not a general requirement for human review of every travel-planning answer (European Commission, 2026a). None of these transparency provisions certifies the truth of a price or itinerary.

The Commission states that Article 50 obligations apply from August 02, 2026, with a transition until December 02, 2026 for the marking and detection obligations concerning systems already on the market before August 02. It also distinguishes provider obligations from professional deployment and personal use (European Commission, 2026b). The identity and activities of the relevant business therefore matter. A company using a third-party model cannot be classified solely from the fact that it has purchased access to that model.

In the UK, the relevant analysis begins with commercial practices rather than a distinct category of AI misinformation. The Digital Markets, Competition and Consumers Act 2024 addresses misleading actions in section 226 and misleading omissions in section 227. Section 229 addresses professional diligence, while section 230 specifies material information for invitations to purchase. Whether a generated statement falls within these provisions depends on the practice and the applicable statutory test. Algorithmic generation does not itself decide the issue (Digital Markets, Competition and Consumers Act 2024).

The CMA explains that the unfair-trading provisions apply to commercial practices from April 06, 2025 and cover traders promoting or supplying their own products or those of other businesses. The general misleading-practice tests involve the likelihood of changing the average consumer's transactional decision, while specified always-unfair practices and invitation-to-purchase information have their own rules (CMA, 2025). A commercial recommendation may therefore warrant scrutiny before a completed booking. These provisions do not, on their own, determine the allocation of every loss between a model supplier, a travel intermediary, and a service provider.

## 5. DISCUSSION

### 5.1. Implications for Tourism Policy

The descriptive and adjusted findings answer different questions. The co-occurrence analysis shows that some consumers who rate AI misinformation potential highly also intend to use AI for travel decisions. The regression analysis shows that, among respondents with similar acceptance-related perceptions, the risk score has a small negative association with intention. These results can coexist because consumers differ in their experience and their expectations of the technology. Treating either result as a direct estimate of the effect of information disclosure would exceed the design.

The policy implication is correspondingly specific. Consumer education may improve understanding, but regulators should not assume that general awareness makes verification and redress unnecessary. Conversely, intended use should not be treated as evidence of irrationality or vulnerability. A tourist can reasonably use a generative assistant for brainstorming while relying on a booking provider or public authority for consequential facts. The regulatory objective should be to make those boundaries visible and usable within the service.

This interpretation is consistent with experimental evidence that errors can influence acceptance through perceived accuracy and trust, with outcomes depending on how the error becomes apparent (Rejón-Guardia et al., 2026). It also recognizes the reputational consequences of failures for tourism businesses (Belanche et al., 2026). The present study translates the coexistence observed in public consumer data into service practices that can be specified, audited, and tested.

### 5.2. Safeguards Across the Travel Decision Process

The first proposed safeguard is a clear account of the service's role. At the start of the interaction, a commercial agent should identify the business offering the service, explain that generative AI is involved where disclosure is required, and state whether it is suggesting options, retrieving current inventory, or authorized to transact. Commercial sponsorship or restricted supplier coverage should be presented when it can materially affect the recommendation. This information should remain accessible as the conversation develops. A single opening notice is unlikely to explain the evidential status of every subsequent claim.

The second safeguard concerns consequential factual assertions. Firms should establish verification procedures for claims capable of causing substantial expenditure, loss of access, or disruption. Examples include total payable prices, cancellation conditions, accessibility features, connection feasibility, and attraction availability. The appropriate source depends on the claim: a current inventory response for a fare, the relevant supplier's conditions for a cancellation term, or an authoritative public source for entry requirements. A link alone is insufficient if it does not support the proposition presented. The system should preserve the relationship between the claim, supporting source, and time checked.

Verification should be proportionate to the task. An inspirational list of possible destinations can tolerate exploratory wording and explicit uncertainty. A statement that a named fare is refundable needs stronger support before the consumer purchases it. Where the system cannot resolve a consequential uncertainty, it should say so, direct the tourist to an appropriate source, and refrain from converting speculation into a transaction instruction. This proposal establishes a performance objective rather than prescribing one retrieval architecture or assuming that current models can eliminate all errors.

The third safeguard is confirmation at the point of commitment. A booking-capable agent should display the selected supplier, dates, total price, principal conditions, and the action about to be taken before expenditure or a material amendment occurs. The consumer's authorization should correspond to the actual action, including its financial limit. An exploratory request to compare hotels should not be interpreted as unlimited permission to purchase. For systems operating within previously delegated limits, the service should make those limits inspectable and permit withdrawal of authority. These proposals address foreseeable consequences of execution; their effectiveness is not established by the 2023 planning survey.

The fourth safeguard is practical correction. A tourist should be able to contest a material claim without recreating the conversation from memory or identifying the technical cause of the mistake. The service should offer an accessible route to a person or

competent review process, preserve a concise record of the disputed recommendation, and correct information that remains in circulation. Where a transaction can still be interrupted or amended, that opportunity should be made available promptly. Whether compensation is owed must be assessed under the relevant law and facts, including the relationship between the representation, reliance, and loss.

These safeguards require an identifiable business responsible for responding to the consumer. As a policy proposal, a trader that markets a branded AI travel service should provide the first point of complaint handling and retain sufficient evidence to investigate its recommendations. Contractual arrangements can then allocate technical investigation and recourse among the trader, model provider, and information supplier. Such arrangements should not require the tourist to diagnose whether an error arose in training, retrieval, summarization, or inventory integration before a complaint can be considered. This is a proposed service obligation, not a claim that existing law imposes universal strict liability on every participant in the supply chain.

### 5.3. Enforcement and Destination Governance

Tourism authorities have a role where commercial advice affects access to destinations and local services. Destination management organizations can maintain current, machine-readable information on public attractions, opening schedules, accessibility, and temporary closures, with clear update responsibility. This would improve the availability of authoritative local information without making the destination organization responsible for every downstream AI summary. Agreements with intermediaries should specify how updates and corrections travel through the information chain. Small operators could contribute through shared templates rather than maintain separate bespoke interfaces for each platform.

Consumer authorities should assess representative travel tasks and preserve evidence of the actual user journey. Relevant indicators include the proportion of material factual claims supported by an appropriate current source, whether unavoidable charges are incorporated at the relevant stage, the frequency of unauthorized actions, and the time needed to correct a verified error. These are proposed evaluation measures; no values for them are estimated in this paper. Testing should include different languages, unfamiliar destinations, accessibility requests, and claims that change shortly before departure. A high average response-quality score could conceal failures concentrated in precisely those circumstances.

Regulatory coordination should also distinguish supervisory functions. AI transparency oversight can examine disclosure and marking, consumer enforcement can examine misleading commercial practices, and tourism bodies can address destination information and sector implementation. These functions may overlap institutionally, but the duties being checked should remain explicit. A disclosure-compliant interface can still display a misleading total price. A factually correct recommendation can still conceal a paid ranking arrangement. Enforcement guidance should explain the conduct at issue rather than assume that an AI label resolves every information problem.

The burden on firms should scale with foreseeable consequences and operational control. Shared verification services, common supplier-data formats, and reusable complaint records could reduce the cost for smaller tourism businesses. Evidence retention should be limited to what is necessary for checking recommendations and handling disputes, with access controls and appropriate deletion practices. Requiring indiscriminate storage of entire conversations could create avoidable privacy risks. A focused record of a material claim, the source used, the authorization given, and the subsequent correction would often be more relevant to accountability.

The proposed approach also needs evaluation of unintended effects. Repeated warnings may become background text, excessive confirmation steps may impede accessible use, and uniform certification costs may favor large intermediaries. A prospective trial should compare a basic AI-origin disclosure with claim-specific support and transaction confirmation, measuring error recognition, verification behavior, decision quality, and complaint outcomes. Consumer satisfaction and intended use should be secondary outcomes where they conflict with factual reliability. This design would test the safeguards directly rather than inferring their success from the observational associations reported here.

## 6. CONCLUSIONS, LIMITATIONS AND FUTURE RESEARCH

Publicly archived consumer data show that stronger perceived hallucination potential can coexist with intended AI use for tourism decisions. Among 224 respondents with high perceived risk scores, 58 endorsed future travel-planning use. The conditional association becomes slightly negative after adjustment for experience and acceptance-related perceptions, but its magnitude and specification sensitivity provide no basis for estimating the effect of a warning or regulation.

The regulatory assessment identifies complementary functions for transparency, substantive commercial accuracy, and practical redress. AI-origin disclosure helps identify the nature of an interaction; it does not validate a travel claim. Existing unfair-trading rules can address misleading commercial conduct, subject to their legal scope and tests. Tourism policy should connect these requirements to proportionate verification of consequential facts, confirmation of authorized commitments, and accessible correction through an identifiable trader. These safeguards should be tested against observed decision quality and recoverability before claims are made about their effectiveness.

The study has several material limitations. The dataset records early consumer responses in 2023 and cannot describe the prevalence of attitudes or the performance of commercial systems in 2026. Recruitment and representativeness cannot be fully assessed from the archived variables. The lack of demographics and country indicators prevents analysis of language, age, legal knowledge, or jurisdiction-specific vulnerability. Confidence intervals quantify uncertainty conditional on the observed analytical design; they do not correct these coverage limitations.

Finally, the legal assessment is limited to selected EU and UK instruments and does not resolve individual disputes. Contract, package-travel, data-protection, accessibility, and sector-specific rules may add obligations depending on the service. The survey contains no warning intervention, error-detection task, verified financial loss, or autonomous transaction. Future research should combine documented agent capabilities with repeated consumer decisions and direct legal-compliance audits. Field studies across source languages and destination types are particularly important before prescribing mandatory sector-wide technical standards.

## REFERENCES


Balaban, D.C., Phua, J. (2026), Is this real or fake? Examining the effects of disclosure of AI-generated content and source type in Instagram-based travel destination advertising on consumer attitudes and behavioral intentions. Computers in Human Behavior, 183, 109047.

Belanche, D., Casaló, L.V., Flavián, C. (2026), When AI fails: The impacts of hallucinations and service errors on customer loyalty and word of mouth. Journal of Travel Research. https://doi.org/10.1177/00472875261456334

Carvalho, I., Ivanov, S. (2024), ChatGPT for tourism: Applications, benefits and risks. Tourism Review, 79(2), 290-303.

Christensen, J., Hansen, J.M., Wilson, P. (2025), Understanding the role and impact of generative artificial intelligence (AI) hallucination within consumers' tourism decision-making processes. Current Issues in Tourism, 28(4), 545-560.

Clarke, R. (2025), Principles for the responsible application of generative AI. Computer Law and Security Review, 57, 106131.

Competition and Markets Authority. (2025), Unfair Commercial Practices (CMA207). Available from: https://www.gov.uk/government/publications/unfair-commercial-practices-cma207/unfair-commercial-practices

Digital Markets, Competition and Consumers Act 2024. Cookies on Legislation (UK). Ch. 13. Available from: https://www.legislation.gov.uk/ukpga/2024/13/contents

Duivenvoorde, B. (2025), Generative AI and the future of marketing: A consumer protection perspective. Computer Law and Security Review, 57, 106141.

European Commission. (2026a), Article 50: Transparency Obligations for Providers and Deployers of Certain AI Systems. AI Act Service Desk. Available from: https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-50

European Commission. (2026b), Transparency Obligations Under Article 50 of the AI Act. Available from: https://digital-strategy.ec.europa.eu/en/faqs/transparency-obligations-under-article-50-ai-act

Hansen, J., Christensen, J., Wilson, P. (2023), Technology Acceptance Model(TAM) and AI Hallucination Survey Dataset. United States: Utah State University.

Huang, G.Q.I., Zhang, C.J., Zhou Torres, W.C., Xiong, X., Li, H. (2026), Responsible AI and human collaboration in tourism management: Ethical considerations and identity disclosure. Tourism Management, 113, 105315.

Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y.J., Madotto, A., Fung, P. (2023), Survey of hallucination in natural language generation. ACM Computing Surveys, 55(12), 248.

Rejón-Guardia, F., Molinillo, S., Anaya-Sánchez, R. (2026), AI hallucinations in tourism: How errors impact consumer trust and recommendation acceptance. Journal of Consumer Behaviour, 25(2), 923-938.

Shen, P., Chen, Y., Xiong, S. (2026), Responsible generative AI in tourism: Conceptualization, measurement, and validation. Journal of Hospitality and Tourism Management, 67, 101441.

Singu, H.B., Chakraborty, D., Troise, C., Camilleri, M.A., Bresciani, S. (2026), Responsible AI for trustworthy tourism: A framework for mitigating ambiguity and anxiety with generative AI. Technological Forecasting and Social Change, 223, 124407.

Sun, D., Wong, I.A., Xiong, X., Li, S. (2026), When cutting edge meets silver tongue: Understanding the word-of-machine effect on travel decisions. Tourism Management, 112, 105271.

Topsakal, Y. (2025), How familiarity, ease of use, usefulness, and trust influence the acceptance of generative artificial intelligence (AI)-assisted travel planning. International Journal of Human-Computer Interaction, 41(15), 9478-9491.