Journal of
Theoretical and Applied Electronic Commerce Research

MDPI

Article

How Competing Platform Actors Frame a Disruptive Technology: A Dyadic Analysis of Generative-AI Risk Disclosure by Hotels and Online Travel Agencies

Yonghee Kim [1], Shi-Shiuan Wang [2] and Sungjin Yoo [3,*]

[1] Division of Business Administration, Sunmoon University, 70 Sunmoon-ro 221beon-gil, Tangjeong-myeon, Asan 31460, Republic of Korea; yhkim1981@sunmoon.ac.kr
[2] Smart Tourism Education Platform (STEP), College of Hotel & Tourism Management, Kyung Hee University, 26 Kyungheedae-ro, Dongdaemun-gu, Seoul 02447, Republic of Korea; jenny8606@gmail.com
[3] College of Business Administration, Soongsil University, 369 Sangdo-ro, Dongjak-gu, Seoul 06978, Republic of Korea
* Correspondence: sjyoo@ssu.ac.kr

Abstract

Hotel executives publicly describe generative AI (GenAI) as their route out of dependence on online travel agencies (OTAs), while OTAs treat the same technology as a threat to be absorbed. The emerging disclosure literature cannot detect this opposition because it pools structurally heterogeneous firms into a single cross-section. This paper asks whether hotels and OTAs frame GenAI differently in their mandatory risk disclosures, and how those frames evolve. This paper analyzes SEC annual-report risk-factor disclosures (FY2019–2025) from a matched dyad of four hotel operators and three OTAs—49 filings, on Form 10-K and, for one foreign private issuer, Form 20-F—combining keyword-anchored extraction, manual thematic coding, cross-firm cosine similarity, and reproducibility-tested stance coding. Three findings emerge. GenAI language is dyadically differentiated in theme: rival-capability framing is exclusive to the OTA side (11 of 51 OTA passages versus none of 24 operator passages), while operator language is dominated by deployment-liability and compliance content (58 versus 28 percent). Firm-specific language nonetheless converges toward common legalistic phrasing within three reporting cycles. And stance is threat-leaning from the moment GenAI enters these filings and intensifies thereafter: threat framing accounts for roughly two-thirds of stance-bearing passages in FY2023 and rises to almost ninety percent by FY2025, while opportunity framing stays flat. The single seven-firm dyad yields illustrative rather than generalizable findings. Yet rivals' annual filings are public, auditable competitive intelligence: OTAs now name supplier-direct and AI-mediated discovery as threats—precisely the opening hotels pursue. A disruptive technology, this evidence suggests, is not merely disclosed but strategically framed from opposing structural positions, and the intensity of its threat framing tracks its competitive salience.

Keywords: generative AI; risk disclosure; online travel agencies; disintermediation; text mining; technology framing; hotel distribution

Check for updates

Academic Editor: Qingfeng Zeng

Received: 16 July 2026
Revised: 14 September 2026
Accepted: 24 September 2026
Published: 26 September 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

When a technology arrives that can relocate the point at which customers discover and choose a product, the two parties to a distribution relationship do not face it symmetrically.

Suppliers stand to recover a channel; intermediaries stand to lose one. Generative artificial intelligence (GenAI) in travel distribution makes that asymmetry unusually legible, because both sides have begun to state their positions on the public record and in opposite terms. In February 2026, the chief executive of a major U.S. hotel franchisor described conversational AI tools to investors as "a unique opportunity" to reduce his company's dependence on online travel agencies (OTAs). The commitment behind the claim was substantial: more than USD 425 million to AI initiatives, and roughly 350 agentic AI agents handling millions of guest calls and reservation requests, which the firm reported were adding "hundreds of basis points" of direct bookings while reducing on-property labor costs for franchisees. Connecting the firm's inventory to a large language model cost under USD 100,000, management noted, and bookings originating in the AI interface carried no per-transaction fee. That figure invites comparison with OTA commissions, which scholarship places at 15–30 percent of room revenue, the burden falling hardest on smaller properties with the least negotiating power [1]. In the same season, the chief financial officer of the world's largest hotel chain observed that AI-originated bookings "could potentially be cheaper than the OTAs". Hotel groups across Asia, Europe and North America have moved in the same direction: expanding loyalty membership, withholding selected inventory from intermediaries, and restructuring their websites so that AI systems can parse them—a practice the trade press has begun to call generative engine optimization.

The intermediaries have not stood still. OTAs have responded by deepening personalization, building recommendation engines that anticipate traveler preferences before a search is entered, and repositioning themselves as content and discovery platforms rather than as booking utilities. Their strategic logic is straightforward: if the point of discovery migrates to an AI assistant, the party that supplies the inventory and the comparison data to that assistant may retain its economic position even as its brand becomes invisible to the traveler. Whether GenAI dissolves the intermediary or merely relocates it is, at the time of writing, genuinely unsettled.

What makes this moment analytically interesting is not that a new technology has arrived, but that the two sides of the same distribution relationship are publicly disagreeing about what it means. Suppliers describe GenAI as their escape route from commission; intermediaries describe it as a threat to be absorbed. The same technology, at the same moment, is being construed in opposite directions by parties whose economic interests are opposed. Trade commentary has noticed the tension. The observation that hotels regard AI simultaneously as the threat and as the remedy recurs across industry coverage of the 2026 earnings season. It has been treated there as an intriguing paradox rather than as an object of systematic study.

This is not the first time the industry has faced such a moment, and the historical precedent counsels caution. Two decades ago, Buhalis and Law [2] forecast that the internet would challenge the role of intermediaries as travelers increasingly transacted directly with suppliers. The forecast proved only half right. The internet did enable direct booking; it also produced a more powerful new intermediary in the OTA [3]. The industry, in other words, has already lived through one cycle in which a technology promising disintermediation delivered re-intermediation instead. That precedent does not tell us how the GenAI cycle will end, but it does establish that the question cannot be answered by reading the technology's affordances alone. It must be answered, at least in part, by attending to how the incumbents themselves interpret the situation—because those interpretations shape what they do.

Here the academic literature falls conspicuously short. Research on hotel distribution has established the economics of the hotel–OTA relationship with considerable precision. It has documented the billboard effect, whereby OTA listing paradoxically raises direct

bookings [4,5]; debated whether OTA participation is financially beneficial on net [6]; and revisited the disintermediation question across successive technological waves [3]. Yet this literature is silent on GenAI, and—more fundamentally—it studies the relationship's outcomes rather than the parties' interpretations of it. Meanwhile, tourism research on AI has grown rapidly but has concentrated almost entirely on the demand side and on internal operations: service robots, automation of the guest encounter, and traveler responses to AI-enabled services [7,8]. It has not asked how AI reconfigures the distribution channel, nor how firms positioned differently within that channel understand what is happening to them.

A third literature has begun to measure firms' exposure to GenAI, but it measures the wrong thing for the present purpose. Eisfeldt et al. [9] construct the first firm-level measure of GenAI exposure and show that equity markets repriced firms sharply in the weeks after ChatGPT's release. Babina et al. [10] measure firm-level AI investment from employee résumés and link it to growth and product innovation. Both are exemplary, and both locate exposure in the firm's workforce—in the tasks its employees perform and in the skills it hires. That is the right construct for asking whether AI substitutes for labor. It is the wrong construct for asking whether AI relocates the customer interface. A hotel operator and an OTA may have near-identical workforce exposure to GenAI while facing diametrically opposed exposures through their positions in the distribution channel. Exposure, in the setting studied here, is positional rather than occupational, and no existing measure captures it.

A fourth and closely related body of work—much of it very recent—counts AI risk language in corporate filings, documenting an abrupt, economy-wide rise in such language after ChatGPT's release and observing that it quickly becomes generic. This work asks how many firms disclose AI risk and how boilerplate that disclosure has become. It is cross-sectional in spirit: it pools structurally heterogeneous firms into a single population and reports a mean. Such a design cannot, by construction, detect the phenomenon described above—two firms in the same value chain, exposed to the same technology, reading it in opposite directions. Pooling a hotel operator and an OTA into one cell averages away precisely the variation that matters. What the industry evidence suggests is a structured opposition; what this literature reports is a population average. The gap is therefore not that AI disclosure is unstudied, but that it has been studied in a way that renders opposition invisible.

This omission matters because framing is consequential, not cosmetic. Entman [11] defines framing as the selection and salience of particular aspects of a perceived reality so as to promote a specific problem definition and evaluation. The strategic-issue literature establishes that such definitions have teeth. Dutton and Jackson [12] show that categorizing an issue as a threat rather than an opportunity systematically alters information processing and subsequent organizational action. Gilbert [13], studying newspaper incumbents confronting digital media, demonstrates that threat perception reshapes both what firms invest in and how rigidly they execute. That setting is a close structural analog to intermediaries confronting GenAI. In platform-mediated distribution, where the intermediary's position rests on operating the interface through which discovery occurs [14,15], suppliers and intermediaries have structurally opposed exposures to any technology capable of relocating that interface. Theory therefore predicts opposition in how they construe GenAI. Nobody has looked.

This study looks at the one archive in which both parties are compelled to speak on the record. Mandatory risk-factor disclosure in SEC annual reports (Item 1A of Form 10-K; Item 3.D of Form 20-F for foreign private issuers) obliges firms to state the material risks facing their business, under legal liability, on an annual cycle. Accounting research

establishes that this text is informative about actual exposures [16] and moves investors' risk perceptions [17], while also warning that it drifts toward generic, boilerplate language that mimics industry peers [17,18]. The methodological toolkit for reading such text at scale—domain-specific dictionaries, topic models, similarity metrics, pattern-based exposure measures—is mature [19–24]. Within tourism, Li et al. [25] have demonstrated that sectoral disclosure repays textual analysis, though their window predates GenAI and their design treats firms as a population rather than as counterparties.

Reading both sides of a distribution relationship against one another, rather than pooling them, is the methodological move that makes the competing-frames question askable. SEC annual-report risk-factor disclosures (FY2019–2025) are assembled from a matched dyad of four major listed hotel operators and three OTAs, and three questions are asked:

RQ1 (differentiation). Do hotel operators and OTAs frame GenAI through different risk themes?

RQ2 (convergence). Does firm-specific GenAI language converge toward common boilerplate over time?

RQ3 (stance dynamics). How does the opportunity-versus-threat balance of GenAI framing evolve across years?

The contribution is conceptual: a disruptive technology is not merely disclosed but strategically framed from opposing structural positions within a distribution dyad, with a threat framing that hardens as competitive salience grows. Section 4 reports what the three questions yield; the design is exploratory and cannot support causal claims about market outcomes, and Section 4.5 states precisely what is and is not tested here.

The remainder of the paper proceeds as follows. Section 2 reviews the literature on hotel distribution, risk disclosure and strategic framing, and locates the research gap. Section 3 describes the corpus, measures and analytic stance. Section 4 presents the findings. Section 5 discusses theoretical and practical implications, together with limitations and directions for future research. Section 6 concludes.

2. Literature Review and Research Gap

2.1. Intermediation and Disintermediation in Hotel Distribution

The economics of hotel distribution turn on a trade-off. OTAs aggregate demand, reduce travelers' search costs and provide reach that individual properties cannot replicate; in exchange, they extract a commission. The theoretical foundation for this exchange is well understood. Bakos [14] showed that electronic marketplaces create value chiefly by lowering buyers' costs of searching for sellers, and that doing so erodes sellers' ability to extract monopolistic profits. His analysis concerns the supplier side; its implication for the intermediary is positional rather than distributive, in that whoever operates the interface through which search occurs occupies the point at which that erosion is administered. Rochet and Tirole [15] formalized the platform's position between two sides of a market, with pricing power derived from its capacity to internalize cross-side network effects. Applied to lodging, these arguments explain why the OTA occupies a structurally advantaged position: it operates the interface through which discovery occurs, and its commission is earned from that position rather than from any operational contribution to the hotel stay itself.

Empirical work in hospitality has examined the consequences of this arrangement from the hotel's side. Anderson [4] documented the "billboard effect", whereby OTA listing raises a hotel's own-channel bookings by exposing the property to travelers who subsequently book direct; Anderson and Han [5] confirmed that the effect persists. This finding complicates the simple story that OTAs merely tax hotels: OTA presence also generates

demand that accrues to the direct channel. O'Connor et al. [6], applying transaction-cost reasoning to U.S. hotel profitability data, find the net financial contribution of OTA participation contested rather than uniformly positive or negative. The strategic question, in short, remains genuinely open.

A parallel analytical literature models the relationship formally, and its results sharpen the intuition. Ling et al. [26] cast the hotel and the OTA as principal and agent: the hotel sets a commission and the OTA, having accepted it, decides how prominently to distribute the hotel's inventory. Guo and He [27] analyze coordination in the tourism supply chain more generally, showing that the division of surplus depends on which party controls the customer-facing decision. The empirical magnitudes are consequential: Toh et al. [1] report commissions of 15–30 percent, borne disproportionately by smaller hotels with the least bargaining power. Rate-parity clauses, which oblige hotels to offer no lower prices on their own websites than on the OTA, have functioned as the enforcement mechanism sustaining this division; their legality has been contested on both sides of the Atlantic, and Nicolau and Sharma [28] examine the consequences of banning them. The picture that emerges is of a relationship in which the intermediary's advantage rests not on superior operational capability but on control of the interface and on contractual terms that protect that control.

This matters for how GenAI should be theorized. If the OTA's position rests on interface control, then a technology that relocates the interface threatens the position directly—not by out-competing the OTA on price or service, but by rendering its interface irrelevant to the traveler's decision. Conversely, if the AI assistant consumes OTA-supplied inventory and comparison data in order to answer travelers' questions, the OTA retains its position while surrendering only its brand visibility. The two scenarios have opposite implications for the OTA's rents, and—crucially—they are not distinguishable from the outside at present. This indeterminacy is what makes the parties' own interpretations informative.

Running through this literature is the recurring question of whether each new technology intermediates or disintermediates the hotel–guest relationship. Buhalis and Law [2], in the landmark review of eTourism, anticipated that the internet would challenge intermediaries as travelers booked directly. Law et al. [3], revisiting disintermediation from the perspectives of both hotels and travel agencies, observed that the outcome was more complicated: the internet enabled direct transactions while simultaneously producing a more powerful new intermediary. The pattern is not confined to travel. Balyuk and Davydenko [29], studying online lending, document re-intermediation in FinTech, where platforms that promised to connect borrowers and lenders directly instead assumed intermediating functions themselves. The historical regularity is therefore re-intermediation rather than disintermediation—a regularity that supplies the natural skeptical prior against which GenAI's disruptive potential should be assessed, and which makes the framing question empirically interesting. If incumbents have seen this film before, do they now read GenAI as a repeat of the last cycle, or as something categorically different?

2.2. Risk-Factor Disclosure as a Window on Managerial Perception

SEC annual reports require firms to disclose the material risks facing their business (Item 1A of Form 10-K; Item 3.D of Form 20-F). A substantial accounting literature establishes that this text is not merely legal boilerplate in every respect: Campbell et al. [16] show that mandatory risk-factor disclosures contain information about firms' actual risk exposures, and Kravet and Muslu [17] show that increases in risk disclosure are associated with higher stock-return volatility, greater trading volume and more dispersed analyst forecasts—that is, that such text shifts investors' risk perceptions rather than resolving uncertainty.

The methodological apparatus for analyzing such text is now mature. Loughran and McDonald [22] demonstrated that general-purpose sentiment dictionaries misclassify financially significant words and constructed a domain-specific alternative; Loughran and McDonald [23] showed that common readability measures are poorly specified for financial disclosure and proposed a simpler proxy; and Loughran and McDonald [24] survey the field. Li [30] classified forward-looking statements using naïve Bayesian methods. Brown and Tucker [20] measured the extent to which firms modify their Management Discussion and Analysis from one year to the next, finding both that firms modify more when underlying economic conditions change more, and that the degree of modification is positively associated with the market's price reaction. Updated narrative disclosure, on their evidence, is informative disclosure. Their modification score is the direct methodological ancestor of the convergence measure employed here, and their finding supplies the benchmark against which convergence should be read. Hassan et al. [21] adapted pattern-based sequence classification from computational linguistics to construct firm-level measures of exposure to an emergent risk (political risk), validating the general strategy of reading corporate text as a barometer of firm-specific exposure. Blei et al. [19] supplied the topic-modeling machinery, and Dyer et al. [18] applied it at scale to three decades of 10-K filings.

Against these constructive findings stands the boilerplate critique. Dyer et al. [18] document that risk-factor text has grown longer, stickier and more redundant over time. Kravet and Muslu [17] report that their results are less pronounced for firm-level disclosures that deviate from those of other firms in the same industry and year, and they read this as lending support to critics' argument that risk disclosures are largely boilerplate. Theirs is the field's canonical statement of the boilerplate concern, though it does not itself estimate the effect of convergence on informativeness. The boilerplate regularity therefore functions in this study as both benchmark and puzzle. If firms converge on shared phrasing even for a genuinely novel risk, then the window during which such disclosure carries firm-specific information may be narrow, and it may close precisely when the underlying competitive situation is diverging most sharply.

2.3. Framing, Threat and the Strategic Interpretation of Technology

Why should the language firms use to describe a technology be of interest beyond its informational content? Framing theory supplies the answer. Entman [11] defines framing as the selection and salience of particular aspects of a perceived reality so as to promote a specific problem definition, causal interpretation and evaluation. On this account, a frame is not a neutral description but an act of construction that shapes what follows.

The strategic-issue literature makes the organizational consequences concrete. Dutton and Jackson [12] argue that categorizing an issue as a threat rather than an opportunity systematically alters both information processing and the motivation of decision makers, and thereby the actions organizations take. The threat-rigidity tradition documents one such consequence: under perceived threat, organizations constrict information processing and centralize control [31]. Gilbert [13], studying newspaper firms confronting the rise of digital media—a structural analog to intermediaries confronting GenAI—shows that the consequences are not uniformly negative: a strong threat perception helped incumbents overcome resource rigidity (they invested) while simultaneously amplifying routine rigidity (they invested in ways that reproduced existing processes).

Two implications follow for the present study. First, how a firm frames a disruptive technology is neither cosmetic nor epiphenomenal; it is an antecedent of organizational action, and therefore worth measuring in its own right. Second, framing may be unstable. If threat-versus-opportunity categorization depends on an issue's perceived controllability and competitive salience, and if those perceptions evolve as a technology matures, then

the frame attached to a given technology should be expected to migrate over time. Prior disclosure research has not examined such migration, in part because it has treated AI risk as a static attribute to be counted rather than as an evolving construction to be traced.

2.4. Artificial Intelligence in Tourism and Hospitality

Research on AI in tourism and hospitality has expanded rapidly. Tussyadiah [8], introducing the Annals of Tourism Research curated collection on AI and robotics, maps a field concentrated on automation of service encounters and its implications for tourist experience. Ivanov and Webster [7] set out an economic research agenda for service robots, spanning tourism supply, demand and destination management. Buhalis and Law [2] remain the reference point for the broader eTourism trajectory within which these developments sit. The dominant orientation of this work is demand- and operations-facing: it asks how travelers respond to AI-enabled services and how firms deploy them internally.

The arrival of large language models has since produced a distinct and fast-growing sectoral literature. Gursoy et al. [32] survey the early conceptual terrain, cataloging the benefits GenAI promises across the stages of travel alongside the threats and challenges it poses to incumbent firms. Dwivedi et al. [33] extend this into a practice-facing agenda, locating implementation challenges from the perspectives of companies, customers and regulators. Dogru et al. [34], drawing on a panel of thirty industry and academic contributors, assemble the broadest framework to date, spanning operations, marketing, revenue management, human resources and strategy. Wong et al. [35] supply the demand-side counterpart: assuming the role of a traveler in a question-and-answer sequence with ChatGPT, they show that the system filters and prioritizes online travel information and returns tailored recommendations—precisely the function that hotel and OTA filings describe, from opposite directions, as either a channel opportunity or a threat to intermediated discovery.

Two features of this work matter for what follows. It is overwhelmingly prospective and normative, asking what GenAI will do to the sector and what firms ought to do in response; and it takes the industry, not the firm, as its unit of analysis. Where threat appears, it appears generically—as a challenge facing hospitality and tourism at large—rather than as a position a particular firm occupies relative to a particular counterparty. Neither the opportunity nor the threat has been read from what firms are themselves obliged to say.

Two adjacent strands are deliberately not rehearsed here, because they are developed at length earlier in this section and repeating them would duplicate rather than extend the review. Algorithm-driven distribution—how intermediaries acquire and monetize control of the point of discovery, and what happens to supplier margins when that control shifts—is treated in Section 2.1, which develops the search-cost and electronic-marketplace literature [14,15] and the hotel–OTA distribution literature specifically [1–6,26–29]. Platform dynamics in tourism and lodging, including two-sided pricing, rate parity and the conditions under which an intermediary's position is reinforced rather than eroded by a new channel technology, are likewise treated in Section 2.1 [15,28]. Readers seeking this paper's engagement with algorithm-driven distribution and platform dynamics should look there; the present section is confined to the generative-AI literature specific to tourism and hospitality, which is where the gap this paper enters is located.

Firm disclosure, by contrast, has received little attention in this sector. Li et al. [25] analyze tourism companies' textual risk exposures in 10-K filings, establishing both the feasibility and the relevance of reading sectoral disclosure as data. Their window, however, predates GenAI, and their design is single-sided: firms are analyzed as a population, not as counterparties within a relationship. To the authors' knowledge, no study examines how the two sides of the hotel–OTA relationship frame GenAI against one another.

2.5. Research Gap

Three bodies of literature bear on the question this paper asks, and each answers part of it while leaving the same part unanswered. The first is distribution economics, reviewed in Section 2.1: it establishes that intermediation is the industry's central strategic fault line, and that prior technological shocks produced re-intermediation rather than disintermediation—but it has not examined how GenAI is being construed by the parties concerned. The second is risk-factor disclosure research, reviewed in Section 2.2: it supplies both the instruments for measuring such construal and the boilerplate benchmark against which it should be assessed—but it pools firms without regard to structural position, and so cannot detect opposition. The third is framing and strategic-issue theory, reviewed in Section 2.3: it predicts that structurally opposed actors should construe the same technology differently, and that such construal should shape action and evolve over time—but this prediction has not been examined in a platform-mediated industry facing GenAI. The present study occupies the intersection of these three bodies of literature, and the method described in Section 3 is chosen to operationalize that intersection.

A note on how this review was assembled. The distribution-economics, disclosure and framing literature were identified from the canonical works in each field and extended by forward and backward citation tracing from them. The generative-AI literature in Section 2.4, which is recent enough that citation tracing is unreliable, was identified by keyword search in Google Scholar, Scopus and Web of Science on combinations of "generative AI", "ChatGPT", "large language model" with "tourism", "hospitality", "hotel" and "online travel agency", restricted to peer-reviewed articles published from 2023 onward. Studies were retained when they addressed generative AI in the tourism or hospitality sector specifically, and each retained study was read against its published abstract rather than cited from a secondary summary.

3. Materials and Methods

3.1. Research Design and Rationale for the Dyadic Case

This study adopts an exploratory, theory-generating design centered on a single supplier–intermediary dyad. This choice is deliberate and warrants defense. The competing-frames question is not answerable in a large, undifferentiated cross-section, because the construct of interest—opposition between structurally distinct positions in a value chain—exists only within a defined relationship. Pooling thousands of firms would average away precisely the variation of interest. The trade-off is that the findings are illustrative rather than statistically generalizable; that trade-off is accepted explicitly and the contribution framed accordingly.

3.2. Corpus

The corpus comprises annual reports filed with the U.S. Securities and Exchange Commission for fiscal years 2019 through 2025 by seven firms: four hotel operators (Marriott International, Hilton Worldwide Holdings, Hyatt Hotels, and InterContinental Hotels Group) and three online travel platforms (Booking Holdings, Expedia Group, and Tripadvisor). Six file Form 10-K, from which the Item 1A risk-factor section is extracted; InterContinental Hotels Group, a foreign private issuer, files Form 20-F, from which the equivalent Item 3.D risk-factor disclosure is extracted. All seven firms have a 31 December fiscal year-end; each FY2025 document is a complete annual report for the year ended 31 December 2025, filed between 10 and 26 February 2026, and none is interim or preliminary; and the panel is balanced at one annual report per firm-year, 49 filings in total—42 on Form 10-K from six firms and seven on Form 20-F from InterContinental Hotels Group.

The seven firms were selected purposively rather than by a mechanical screen: the construct of interest exists only within a defined supplier–intermediary relationship, so the design requires a matched dyad rather than a sample drawn to represent a population. Large, listed, pure-play firms were selected on each side of the hotel–OTA relationship with continuous SEC annual-filing coverage across the window; no market-capitalization threshold or top-*N* rule was applied, and no claim is made that the set exhausts or statistically represents its categories.

Extraction proceeds in two stages. The 49 risk-factor sections yield 19,721 sentences. The 290 containing an explicit artificial-intelligence anchor are retained, and within these a strict generative-AI subset is identified, of 75 sentences containing at least one of: "generative AI", "generative artificial intelligence", "gen AI", "large language model", "LLM", "agentic", or "AI agent(s)". The pattern is word-boundary-anchored and Unicode-normalized, is unit-tested against known false positives (including the "LL.M." law degree), and each retained sentence records its matching anchor(s) in the replication files (Appendix A). Because firms repeat risk language across years, the 75 passages comprise 66 distinct texts. One compositional fact conditions all side-level comparisons: Booking Holdings—whose house style abbreviates the technology as "Gen AI"—contributes 42 of the 51 OTA passages; firm-level tables are therefore reported alongside every side-level quantity. The window brackets the public release of ChatGPT in late 2022, permitting a pre- and post-shock view; FY2019–2022 is treated as the pre-GenAI baseline and FY2023–2025 as the GenAI period.

3.3. Measures

Following established practice in disclosure text analysis [21,24], the design combines transparent, keyword-anchored extraction with manual thematic coding, similarity measurement, and model-based stance coding, using unsupervised topic modeling only as exploratory corroboration. Four measures are constructed.

GenAI mention intensity. GenAI-anchored terms are counted—"generative AI", "large language model", "agentic" and cognates—and normalized by document length, yielding a firm-year measure of disclosure intensity. This measure establishes onset and growth, and provides context rather than contribution.

Thematic composition (RQ1). Each of the 75 strict GenAI passages is assigned one of five themes under a written codebook: (i) AI-mediated discovery and channel disintermediation—GenAI relocating travel search, discovery, or booking away from the firm's channel; (ii) rival capability asymmetry—competitors or large technology firms using GenAI to compete more effectively, without the channel-relocation mechanism; (iii) own deployment and adoption; (iv) deployment liabilities and compliance, including AI-enabled fraud and security exposure; and (v) general or descriptive mentions. A model-assisted first pass proposed each code; an independent second stage of three further model passes, run under randomized passage orders and combined by modal label, re-coded all 75 passages (agreement with the first pass 84.0 percent, Cohen's $\kappa = 0.79$), and each of the twelve disagreements was adjudicated against the codebook, three codes changing as a result, with the full adjudication log in the replication files. Identical sentences repeated across years carry identical codes by construction. Passage counts by side are reported together with equal-firm-weighted shares and the full firm-level panel, because one firm dominates the OTA side of the corpus. As a machine-learning complement, latent Dirichlet allocation topic models were also estimated; with 75 short documents, their document–topic prevalences are not stable across random initializations (Appendix B reports the $K = 2$–12 validation sweep and 50-seed sensitivity ranges), so they are treated as exploratory corroboration rather than evidence.

Boilerplate convergence (RQ2). For each fiscal year, pairwise cross-firm cosine similarity of AI-related language (the 290-sentence anchored corpus) is computed and averaged, and reported both raw and under a fixed-content control that equalizes the number of sentences per firm-year. Rising similarity indicates convergence toward shared phrasing. This operationalizes the boilerplate construct of Kravet and Muslu [17] and adapts the year-over-year modification logic of Brown and Tucker [20] from the within-firm to the cross-firm dimension: where Brown and Tucker ask whether a firm updates its own narrative relative to its prior filing, the present measure asks whether firms' narratives become indistinguishable from one another.

Stance (RQ3). GenAI passages are coded as opportunity-oriented ("offensive") or threat-oriented, following the threat–opportunity categorization of Dutton and Jackson [12]. Their construct is explicitly evaluative: an issue is framed as an opportunity when it is construed as controllable and as carrying potential gain, and as a threat when construed as uncontrollable and as carrying potential loss. This is operationalized in the disclosure context as follows. A passage is coded offensive when GenAI is presented as a capability the firm deploys, invests in or benefits from—typically signaled by the firm as grammatical agent ("we are integrating", "we expect efficiencies"). A passage is coded threat when GenAI is presented as a development acting upon the firm, its channel or its competitive position—typically signaled by the firm as grammatical object or by explicit adverse-outcome language ("may reduce", "could result in loss of"). Passages that are purely regulatory or purely descriptive are coded neutral and reported as a separate category; stance shares are computed over the stance-bearing (offensive or threat) passages, and every table states that denominator.

Stance was coded for all 75 passages—the full extracted set, not a subsample—by a locally hosted large language model (Qwen2.5 14B Instruct, model tag qwen2.5:14b-instruct, served through Ollama, https://ollama.com, accessed on 23 August 2026, at temperature 0.2 and top-p 0.9) under the written rule above, five times, each pass with a different random seed (101, 211, 307, 401 and 503) and an independently randomized passage order, in chunks of 25 with a missing-identifier retry. Because seed, order, and batch composition vary together, the agreement is reported across passes as multi-seed, multi-order test–retest stability: mean pairwise Cohen's $\kappa = 0.68$ (range 0.62–0.78 across the ten pass pairs), mean agreement 82.9 percent (range 78.7–88.0), with 50 of 75 passages labeled identically on all five passes; a further one-passage-at-a-time pass agreed with the five-pass modal label on 90.7 percent of passages. The 14 passages on which the passes split three-to-two were re-examined individually against the rule and adjudicated, with the log of votes, final labels, and rationales submitted as Supplementary Table S3; results are reported under both the modal and the adjudicated labels. Test–retest stability establishes that the labels follow reproducibly from the stated rule; it is not evidence of construct validity, which rests on the rule itself and on the adjudication of contested cases. Given the small set, counts are reported rather than inferential statistics, and the measure is treated as exploratory throughout. The exact prompts, seeds and passage orders are provided in the Supplementary Materials; the raw model responses are available from the corresponding author on request.

3.4. Analytic Stance and Scope Conditions

The unit of analysis is the firm-year risk-factor text; inferences are descriptive and interpretive. No causal claims are advanced that GenAI altered distribution outcomes. Supplementary market-based tests were attempted—event studies around AI product launches, difference-in-differences on OTA take-rate and synthetic control—and are reported in Appendix E. The number of firms is small, the GenAI shock arrives contemporaneously

for all of them (removing any clean control group), and event windows are short. These designs are therefore underpowered and, in the difference-in-differences case, formally unidentified. They are reported transparently rather than omitted, in order to forestall the misreading that a null result of this kind constitutes evidence of no effect.

Use of generative AI in the analysis. In accordance with the journal's disclosure requirements, this section states how generative AI was used in producing the results reported here. Thematic coding of the 75 strict GenAI passages used a model-assisted first pass and an independent second stage of three further model passes under randomized passage orders, combined by modal label; stance coding of the same 75 passages was performed five times by a locally hosted Qwen2.5 14B Instruct model, served through Ollama (https://ollama.com, accessed on 23 August 2026) at temperature 0.2 and top-p 0.9, under the written rule of this section, with random seeds 101, 211, 307, 401 and 503 and independently randomized passage orders, followed by a sixth one-passage-at-a-time pass at seed 601. The thematic passes used the same model under the same decoding settings. In both cases, the model proposed labels and the authors adjudicated every disagreement against the written codebook; the authors have reviewed and edited the output and take full responsibility for the content of this publication. Corpus construction, similarity computation, and topic modeling used Python 3 with scikit-learn and gensim (https://www.python.org, https://scikit-learn.org and https://radimrehurek.com/gensim, all accessed on 23 August 2026). The exact prompts, seeds and passage orders are provided in the Supplementary Materials; the raw model responses are available from the corresponding author on request.

4. Results

4.1. Onset of GenAI Disclosure

GenAI-anchored language is essentially absent through FY2022 and appears abruptly across both sides of the dyad in FY2023, growing thereafter. This mirrors the economy-wide onset documented elsewhere and is reported as context rather than as contribution: it establishes the shared exogenous shock against which dyadic differences can be read. Notably, onset is simultaneous across the dyad; whatever differentiates operators from OTAs, it is not the timing of their attention to the technology.

4.2. Dyadic Differentiation of GenAI Themes (RQ1)

Thematic coding reveals a differentiated structure, reported in Table 1. Two contrasts are robust to how firms are weighted. First, rival-capability framing—competitors or large technology firms wielding GenAI—appears in eleven OTA passages and in none of the operator passages; operators never once name rivals' GenAI capability as a risk. Second, the composition of operator GenAI language is dominated by deployment-liability and compliance content: 58 percent of operator passages, against 28 percent of OTA passages, and the operator share rises to 65 percent under equal-firm weighting. Discovery-and-disintermediation framing—GenAI relocating travel search and booking away from the firm's channel—is more frequent on the OTA side in raw counts (21.6 versus 12.5 percent). Every one of its eleven OTA passages comes from Booking Holdings, however, so the pattern is described here as Booking-led rather than as an OTA-sector property. Notably, by FY2025 each of the three operators also names AI platforms as a potential diversion of bookings from its direct channels: the discovery threat has become two-sided.

Table 1. Thematic composition of strict generative AI (GenAI) passages, by side of the dyad. OTA = online travel agency.

Theme	OTA Passages (Share)	Operator Passages (Share)	Equal-Firm OTA/Operator
AI-mediated discovery and disintermediation	11 (21.6%)	3 (12.5%)	0.09/0.11
Rival capability asymmetry	11 (21.6%)	0 (0%)	0.28/0.00
Own deployment and adoption	12 (23.5%)	3 (12.5%)	0.45/0.13
Deployment liabilities and compliance	14 (27.5%)	14 (58.3%)	0.16/0.65
General/descriptive	3 (5.9%)	4 (16.7%)	0.02/0.13

Notes: $N = 75$ strict GenAI passages (51 OTA, 24 operator; 66 distinct texts) from 49 annual-report risk-factor sections, FY2019–2025, coded under the five-category codebook of Section 3.3. Booking Holdings contributes 42 of the 51 OTA passages; the equal-firm columns weight each firm identically and are rounded to two decimals, so they need not sum to one, and the full firm-level panel is in Appendix D. Because the strict subset first appears in FY2023, all passages date from FY2023–2025.

A methodological concern arises immediately: could the differentiation be an artifact of overall verbosity, with one side simply saying more about GenAI in general? Composition, not volume, is what Table 1 measures—each side's shares sum to one—so the contrast cannot be produced by one side's greater output alone. Its validity is corroborated by an independent measurement: the stance codes of Section 4.4 align with the themes exactly as theory predicts, with deployment-themed passages carrying all eight opportunity codes and only two of the forty-seven threat codes, while discovery- and rival-themed passages are coded as threat twenty-two times and opportunity-oriented never. A stronger claim would not be supported by the data. Hotel operators and OTAs are structurally different businesses whose risk disclosures differ on many non-AI dimensions as well (Appendix C), and no claim is made here that GenAI is the only topic on which they diverge.

That caution can be made precise. Figure 1 places the strict generative-AI share alongside three non-AI benchmark categories on a single scale: the share of a firm-year's risk-factor sentences over FY2023–2025. On that scale, the two sides differ on generative AI (0.67 versus 1.32 percent, $|d| = 0.51$). They also differ, by as much or more, on cybersecurity and data privacy ($|d| = 0.68$), on macroeconomic and travel-demand conditions ($|d| = 0.97$) and on workforce and labor ($|d| = 1.88$). This is reported rather than suppressed: hotel operators and online travel platforms are structurally different businesses that write differently about most of their risks, and a demonstration that they differ on generative AI and on nothing else was never available. The verbosity concern is therefore answered not by cross-category flatness but by within-corpus composition, which is what Table 1 measures.

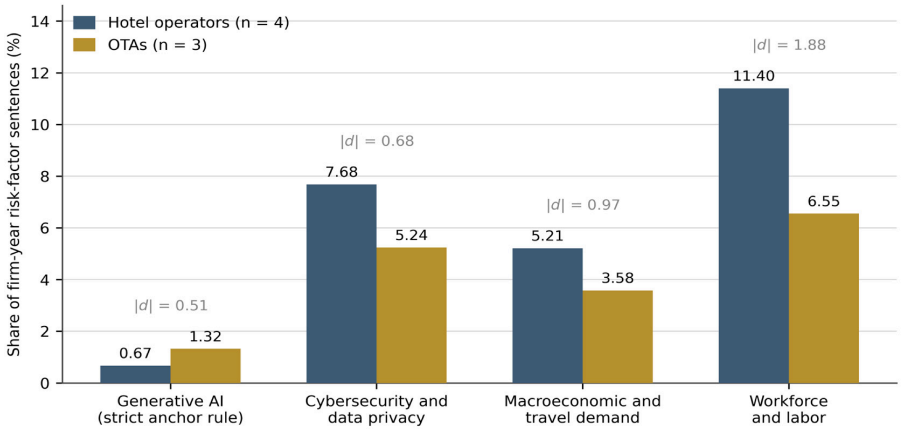


Figure 1. Hotel–OTA differences in risk-sentence composition, FY2023–2025, for generative AI and three non-AI benchmark categories. Each bar is the mean share of a firm-year's risk-factor sentences

matching the category dictionary; $|d|$ is the absolute standardized difference between sides. Non-AI baselines are not flat across the dyad: each differs by as much as or more than generative-AI incidence does. The underlying values appear in Table A5; the category dictionaries and the construction script are available from the corresponding author on request. OTA = online travel agency.

4.3. Convergence Toward Boilerplate (RQ2)

Mean cross-firm cosine similarity of AI-related risk language rises monotonically across the GenAI period: 0.081 in FY2023, 0.125 in FY2024, and 0.212 in FY2025 (Table 2), with six firms contributing AI language in FY2023 and all seven thereafter. Part of this rise reflects volume rather than convergence. The average number of AI-related sentences per firm-year more than doubles over the window, from 8.2 to 19.0. Under a fixed-content control that samples an equal number of sentences per firm-year, the series is 0.086, 0.118 and 0.117. Language converges most sharply between FY2023 and FY2024; thereafter, firms mainly say more, in already-similar terms. Both rows are reported in Table 2, and the full seven-year series—which begins higher in FY2019, when very few AI passages existed per firm, and falls before rising—is in Appendix C.

Table 2. Convergence of AI-related language toward common phrasing.

Fiscal Year	Mean Cross-Firm Cosine Similarity	Fixed-Content Control	Firms
FY2023	0.081	0.086	6
FY2024	0.125	0.118	7
FY2025	0.212	0.117	7

Notes: Mean pairwise cross-firm cosine similarity of AI-related risk-factor text (290-sentence anchored corpus), by fiscal year, seven-firm dyad. The fixed-content control equalizes the number of sentences per firm-year (three, averaged over five random draws), separating convergence of language from growth in its volume.

This pattern is consistent with the boilerplate regularity documented by Kravet and Muslu [17] and with the growing stickiness and redundancy of risk-factor text reported by Dyer et al. [18]. Its speed, however, is notable. Prior work established that risk disclosure tends toward genericness. The present data suggest that this tendency asserts itself very quickly, even for a genuinely novel technology, and even between firms whose actual exposures to that technology are sharply divergent, as Section 4.2 demonstrates.

The implication is pointed, though it must be stated with care. Brown and Tucker [20] find that firms modify their narrative disclosures more when underlying economic conditions change more, and that the extent of modification is positively associated with the market's price reaction: updated language, on their evidence, is informative language. Read against that benchmark, convergence toward a shared template is the opposite of updating. Firms' AI disclosures increasingly resemble one another even as their actual exposures, documented in Section 4.2, diverge. The informational consequence cannot be estimated directly here, and no such estimate is claimed. The conjunction of converging language and diverging exposure is nonetheless difficult to reconcile with disclosure that tracks firm-specific circumstances, and it is consistent with the boilerplate concern that Kravet and Muslu [17] articulate.

4.4. Stance Intensification Across the GenAI Period (RQ3)

Stance coding of all 75 strict GenAI passages shows intensification rather than reversal (Table 3). Threat framing already outweighs opportunity framing among stance-bearing passages when GenAI first appears: in FY2023 the offensive-to-threat count is 2:4 under the modal labels, and 2:5 under the adjudicated labels. Its dominance then grows to 3:18 in FY2024 and 3:25 in FY2025, a threat share rising from roughly two-thirds to ninety

percent. Opportunity framing never exceeds three passages in any year. Volume moves with salience: 10, 26, and 39 GenAI passages in the three fiscal years, with most of the FY2024–2025 growth contributed by Booking Holdings, and neutral or descriptive passages holding roughly a third of the corpus throughout. What changes across the window is therefore not the sign of the evaluation but its dominance and its volume.

Table 3. Stance composition across the GenAI period.

Fiscal Year	Offensive-to-Threat (Modal)	Adjudicated	Neutral	Threat Share of Stance-Bearing
FY2023	2:4	2:5	4 (3)	67% (71%)
FY2024	3:18	3:16	5 (7)	86% (84%)
FY2025	3:25	3:24	11 (12)	89% (89%)

Notes: All 75 strict GenAI passages, coded as offensive, threat, or neutral; modal label across five test–retest passes (mean pairwise $\kappa = 0.68$), with author-adjudicated values in parentheses for the 14 three-to-two splits. Threat share is computed over stance-bearing (offensive + threat) passages only. Counts are reported rather than inferential statistics; the measure is exploratory.

This dynamic matters theoretically. Dutton and Jackson [12] treat threat-versus-opportunity categorization as an antecedent of organizational action, and Gilbert [13] shows that threat perception reshapes both resource and routine rigidity in incumbents confronting a discontinuity. Both bodies of literature take the categorization as an input and trace its consequences. The present data suggest that the categorization is itself an output, but one that hardens rather than flips: a technology that arrives already classified as a threat becomes progressively more so as its competitive salience grows. That the balance moves monotonically, for the same technology and the same firms, indicates that framing intensity is a moving quantity rather than a fixed property of an issue.

Two Interpretations of the Stance Shift

Two readings are available, and the present data cannot fully adjudicate between them. The intensification may reflect genuine strategic escalation: firms recognizing, year by year, that GenAI's competitive implications are larger than first assessed. Alternatively, it may reflect legal caution: as litigation risk around AI disclosure became salient, firms may have added threat language defensively without any change in underlying belief. The boilerplate convergence documented in Section 4.3 lends some support to the second reading, since defensive language is precisely what converges. The absence of any opportunity-dominant early phase strengthens that reading further, as does the concentration of operator-side threat language in deployment-liability boilerplate (Table 1). This is flagged rather than resolved. Distinguishing the two readings would require data on internal deliberation that public filings cannot supply. Given the small coded set ($N = 75$, the full strict subset), the stance result is presented as an exploratory pattern warranting replication rather than as an estimated effect.

4.5. Scope of Inference: Why Market Effects Are Not Tested Here

A reader may reasonably ask whether intensifying threat language corresponds to any market reaction. This study does not test that question, and the supplementary specifications attempted here cannot answer it. Event-study, difference-in-differences, and synthetic-control designs were estimated and are reported in full in Appendix E. None is interpretable. The difference-in-differences specification is not identified: the arrival of generative AI is a shock common to all seven firms, so no clean control group exists. The event-study and synthetic-control estimates rest on seven firms and short windows and are underpowered by construction. Their null results are therefore uninformative—they neither support nor exclude a market effect—and no descriptive claim is drawn from them

anywhere in this paper. They are reported so that readers see what was attempted rather than only what survived, and to specify what a better-identified design would require: staggered, firm-specific adoption events and a substantially larger cross-section. Section 5.4 develops that agenda.

5. Discussion

5.1. Theoretical Implications

The findings extend disclosure and framing scholarship in three ways.

First, they show that a disruptive technology is framed dyadically. Intermediaries and suppliers within the same distribution relationship construe the identical technology from opposed positions—OTAs through the competitive and channel exposure it creates, operators through the liabilities of deploying it. This is consistent with framing theory [11] and with the strategic-issue tradition [12], but it has not previously been demonstrated for GenAI in a platform-mediated industry. The finding complements and qualifies the large-sample AI-disclosure literature: by pooling structurally heterogeneous firms, that literature cannot by construction observe structured opposition, and so risks reporting an average that describes no one. This result suggests that AI-disclosure research should condition on position in the value chain, not merely on industry codes.

Second, they document boilerplate convergence for a genuinely novel risk, extending the boilerplate concern articulated by Kravet and Muslu [17] and the trends documented by Dyer et al. [18] from established risk categories to an emergent one, and adapting Brown and Tucker's [20] modification logic to the cross-firm dimension. The speed of convergence—visible within three reporting cycles, and concentrated in the first of them once volume growth is controlled—suggests that the window during which AI risk disclosure carries firm-specific information may be short. This has a direct bearing on the ongoing debate about disclosure quality: if novel risks are absorbed into shared templates almost immediately, then mandatory disclosure may be structurally ill-suited to communicating precisely those risks that are most novel.

Third, they surface stance dynamics. The monotone intensification of threat framing indicates that evaluative categorization is competitively contingent and moves in intensity even when it does not reverse in sign. Prior work treats threat-versus-opportunity labeling as an antecedent [12,13,31]; the present results show that the label's dominance is itself a moving object, and locate that movement in firms' own mandated disclosures rather than in interviews or surveys. This offers a scalable, archival window on strategic issue diagnosis—a construct that has historically been difficult to observe at scale—and suggests that the disclosure archive may be a productive site for studying managerial cognition longitudinally.

5.2. Practical Implications

For hospitality managers, the dyadic pattern constitutes a strategic signal that is public, auditable and updated annually. OTAs' own filings now name supplier-direct and AI-mediated discovery as threats to intermediation—which is precisely the opening that hotels pursue when investing in direct channels and AI visibility. Systematic monitoring of partners' and rivals' disclosure language is therefore a low-cost competitive-intelligence practice, requiring no proprietary data and no privileged access. That an intermediary is willing to state, in a document carrying legal liability, that its position is threatened, is a stronger signal than any claim it might make in a commercial negotiation.

For investors and analysts, the convergence result is a caution. Once AI risk language becomes universal, its mere presence in a filing carries little diagnostic value; what remains potentially informative is the residual—the firm-specific content that departs from the

shared template. Brown and Tucker [20] show that markets respond to narrative disclosure that is genuinely updated. The results suggest that, for GenAI, the proportion of such content is falling, and that analysts should therefore attend to what a firm says that its peers do not.

For regulators concerned with disclosure quality, the combination of near-immediate threat classification and rapid boilerplate formation suggests that current risk-factor practice may generalize fast-moving technological risks into standardized language almost as soon as it recognizes them. This is directly relevant to ongoing scrutiny of the adequacy of AI-related disclosure. A disclosure regime that produces universal, undifferentiated AI risk language across firms with sharply divergent exposures is not obviously serving investors.

5.3. Boundary Conditions and Generalization

Sections 5.3 and 5.4 answer different questions and are kept separate. This section asks a theoretical question—where the dyadic pattern should be expected to recur, and where it should not—and derives predictions that other studies can test. Section 5.4 records the empirical limitations of the present study, including its single-dyad design; those limitations are stated there and are not restated here.

How far does the dyadic pattern travel? The theory developed here suggests two boundary conditions, both of which are testable in other settings. The first is interface dependence: dyadic opposition should be sharpest where the intermediary's rents derive from controlling the point of discovery rather than from physical logistics or risk-bearing. Lodging is close to the polar case, since the OTA neither holds inventory nor operates the property; it sells access to attention. Where an intermediary performs substantive operational functions—warehousing, fulfillment, underwriting—a technology that relocates discovery threatens less of its value, and correspondingly muted opposition would be expected in its disclosure. The second condition is substitutability of the supplier's own channel: hotels can, in principle, sell direct, which is what makes AI-mediated discovery an opportunity for them rather than merely a new tax. Suppliers lacking a viable direct channel have no such upside and should frame the same technology defensively, alongside the intermediary rather than against it.

These conditions generate concrete predictions, ordered by the degree to which the intermediary's value derives from discovery rather than from operational function. Restaurants and food-delivery platforms should exhibit opposition comparable to lodging, since the platform's contribution is substantially discovery and the restaurant retains a viable direct channel. Marketplaces that also perform fulfillment—warehousing, logistics, payment guarantees—should exhibit weaker opposition, because a technology that relocates discovery threatens a smaller share of what they do. Ride-hailing, where the platform additionally performs matching and pricing under real-time constraints, should be weaker still. And within lodging, independent properties—which depend on intermediaries for discovery and lack loyalty programs—should frame GenAI less optimistically than the branded chains observed here, whose loyalty assets give them a credible direct alternative. The present sample cannot test these predictions, but it renders them precise, which is the proper contribution of a theory-generating study.

5.4. Limitations and Future Research

Five limitations bound this study. First, the sample is a single distribution dyad of seven firms; findings are illustrative and theory-generating rather than statistically generalizable, and they are framed accordingly throughout. Second, the evidence is textual: disclosure framing cannot be linked here to realized channel outcomes, because property-level channel-mix data are proprietary and typically accessible only through firm partner-

ships [4,6]. Third, causal identification of GenAI's market effects is infeasible with these data, as Appendix E documents, owing to the common-shock timing and the small number of firms. Fourth, stance and theme coding rest on a small set ($N = 75$, of which 66 are distinct texts) and on model-assisted procedures validated by test–retest stability, an independent second coding stage, and logged adjudication rather than by independent double-blind human coding; the intensification of threat framing should be treated as a pattern to be replicated rather than as an estimate. Fifth, the corpus is concentrated: Booking Holdings contributes 42 of the 51 OTA passages, so pooled OTA quantities primarily describe that firm, and the text flags throughout which contrasts survive equal-firm weighting.

Three avenues for future research follow directly. The first is to test the boundary conditions derived in Section 5.3, which specify the settings in which the dyadic pattern should sharpen and those in which it should weaken. Section 5.3 states the predictions; what is required here is comparative evidence across those settings to adjudicate them. The second is to connect framing to outcomes. Dutton and Jackson [12] and Gilbert [13] both predict that threat framing changes what firms do; testing this in the present setting requires property- or transaction-level channel data, which would permit examination of whether firms that frame GenAI as an opportunity subsequently shift channel mix. The third is to recover identification by exploiting staggered, firm-specific AI events rather than the common ChatGPT shock, thereby constructing designs with genuine control groups—the absence of which is the binding constraint on causal inference here.

6. Conclusions

Generative AI has entered the risk disclosures of hotels and OTAs alike, but not in the same voice. Reading both sides of the hotel–OTA distribution relationship together, the analysis indicates that GenAI is framed dyadically—through competitive and channel exposure by intermediaries, and through deployment liabilities by suppliers—that firm-specific language converges toward boilerplate within three reporting cycles, and that evaluative stance, threat-leaning from the moment the technology enters these filings, hardens further as competitive salience rises.

The three research questions can now be answered directly. On RQ1, generative-AI language is dyadically differentiated in theme rather than merely in volume: rival-capability framing occurs exclusively on the intermediary side, appearing in 11 of 51 OTA passages and in none of the 24 operator passages, while operator language is dominated by deployment-liability and compliance content, 58 percent against 28 percent. On RQ2, that differentiation does not persist as distinctive prose. At fixed content volume, cross-firm similarity rises across FY2023–2025, so firms converge on common legalistic phrasing within roughly three reporting cycles even as what they face diverges. On RQ3, evaluative stance is threat-leaning from the moment the technology enters these filings and hardens thereafter, from roughly two-thirds of stance-bearing passages in FY2023 to almost ninety percent by FY2025, while opportunity framing stays flat. Read together, the convergence result bears on how the stance result should be interpreted: language that is simultaneously hardening and homogenizing is more consistent with defensive legal drafting than with seven firms independently arriving at the same strategic re-appraisal.

The contribution follows from reading a distribution relationship rather than a population of firms. Theoretically, the paper shows that mandatory risk disclosure carries a structural signature: the same technology is construed as opportunity or as threat according to the position a firm occupies in the channel, which makes disclosure a source of evidence on strategic interpretation and not only on risk quantity. It also shows that this signature is perishable, since the convergence toward boilerplate erodes precisely the firm-specific content that makes framing legible. Practically, the implication is immediate: rivals' annual

filings are public, auditable competitive intelligence, and they are most informative in the first two or three years after a technology arrives, before standardized language closes over the differences. Hotel groups reading OTA filings in that window find their counterparties naming supplier-direct booking and AI-mediated discovery as threats—the precise opening those groups are pursuing.

The broader takeaway is that, in platform-mediated industries, disclosure is not a neutral mirror held up to a technology but a contested construction of it, assembled from opposing structural positions and hardening over time. The industry has been here before—Buhalis and Law [2] expected the internet to disintermediate travel and it produced a more powerful intermediary instead—and whether generative AI repeats that pattern is not yet knowable. What is knowable is that the parties themselves do not agree on the answer, and that their disagreement is legible in the public record for anyone who reads both sides of the relationship at once.

Supplementary Materials: The following supporting information can be downloaded at: https://www.mdpi.com/article/10.3390/jtaer21100339/s1, Table S1: Anchor list, regular expressions, and unit-test summary for corpus construction; Table S2: Thematic coding record for the 75 strict GenAI passages; Table S3: Adjudication log recording, for each contested passage, the pass votes, the final label and the rationale; File S1: Exact model prompts, random seeds and passage orders used in the coding passes.

Author Contributions: Conceptualization, S.Y.; methodology, Y.K.; software, Y.K.; validation, Y.K. and S.-S.W.; formal analysis, Y.K.; investigation, Y.K. and S.-S.W.; resources, S.-S.W.; data curation, Y.K.; writing—original draft preparation, Y.K. and S.-S.W.; writing—review and editing, S.Y. and S.-S.W.; visualization, Y.K.; supervision, S.Y.; project administration, S.Y. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable. The study analyzes publicly filed corporate disclosure documents and involves no human participants.

Informed Consent Statement: Not applicable.

Data Availability Statement: The annual report filings analyzed in this study are in the public domain and can be retrieved from the U.S. Securities and Exchange Commission EDGAR system (https://www.sec.gov/edgar, accessed on 23 August 2026); the seven issuers and the fiscal years are specified in Section 3.2 and Appendix A, so the source documents can be assembled independently. The coding instruments needed to reproduce the analysis—the anchor list and regular expressions, the thematic coding record, the adjudication log, and the exact model prompts, seeds and passage orders—are provided in the Supplementary Materials; the thematic codebook and the stance-coding rule are stated in full in Section 3.3 and Appendix D. The derived sentence corpus, the raw model responses and the analysis scripts are not publicly deposited, because they reproduce extended passages of the source filings and were assembled through a bespoke extraction pipeline; they are available from the corresponding author on reasonable request.

Acknowledgments: During the preparation of this manuscript, the authors used large language models for the thematic and stance coding described in Section 3, where the models, versions, decoding parameters and random seeds are specified. The authors have reviewed and edited the output and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Corpus Construction and Provenance

The extraction rule retains a sentence as strict GenAI when it contains at least one of seven anchors: “generative AI”; “generative artificial intelligence”; “gen AI” (with

or without hyphen, including the one-word form); "large language model(s)"; "LLM(s)" (undotted form only—the dotted form is indistinguishable from the LL.M. law degree, a documented false positive); "agentic"; or "AI agent(s)". All patterns are word-boundary-anchored; Unicode dashes and non-breaking spaces are normalized before matching; and the rule passes a unit-test suite of 13 positive and 9 negative cases (including "cash-generative", "regenerative", and law-degree constructions). Each retained sentence records its matching anchor(s), the rule version, and the source-file hash in the replication files. Table A1 reports the corpus flow and the firm composition of the strict subset.

Table A1. Corpus flow and firm composition.

Stage		***N***
Annual-report risk-factor sections (7 firms × FY2019–2025)		49
Risk sentences		19,721
Sentences with an explicit AI anchor		290
Strict GenAI subset		75
Distinct texts within the strict subset		66
Firm	**Side**	**Strict GenAI passages**
Booking Holdings	OTA	42
Expedia Group	OTA	6
Tripadvisor	OTA	3
InterContinental Hotels Group	Operator	9
Hilton Worldwide Holdings	Operator	7
Hyatt Hotels	Operator	6
Marriott International	Operator	2

Notes: The upper panel reports the corpus flow; the lower panel, whose column headers are set in bold, reports the firm composition of the strict generative AI (GenAI) subset. OTA = online travel agency.

Anchor incidence within the 75 passages: "gen AI" 39; "generative AI" 18; "generative artificial intelligence" 12; "agentic" 8; "large language model" 4; "AI agent(s)" 4; "LLM" 2 (a sentence may match several anchors). One passage is retained solely by the "AI agent(s)" anchor. No strict GenAI sentence occurs in FY2019–FY2022 in any of the 28 pre-period firm-years.

Appendix B. Topic-Model Sensitivity Analysis

Table A2 reports the validation sweep and Table A3 the fifty-seed sensitivity analysis. Latent Dirichlet allocation was estimated on the 75-sentence strict corpus (scikit-learn, batch variational Bayes, max_iter = 50, symmetric priors $\alpha = \eta = 1/K$; CountVectorizer, English stop words, min_df = 2, max_df = 0.5, unigrams and bigrams; text lower-cased with non-alphabetic characters stripped, no lemmatization; coherence from gensim, c_v, twelve top terms per topic).

Table A2. Validation sweep, K = 2–12 (random_state = 42).

K	**c_v Coherence**	**Training Perplexity**
2	0.245	644.7
3	0.521	623.6
4	0.605	583.8
5	0.665	583.6
6	0.711	568.5
7	0.660	544.0
8	0.671	549.1

Table A2. *Cont.*

K	c_v Coherence	Training Perplexity
9	0.633	566.5
10	0.518	548.4
11	0.579	571.4
12	0.607	551.9

Coherence peaks at $K = 6$ on this seed, but no single K is defensible as evidence at this corpus size. To quantify algorithmic stability, each K from 2 to 6 was refit across 50 random initializations, aligning topics to the seed-42 solution by Hungarian matching on term distributions, and the median OTA–operator prevalence gap is reported with its 5th–95th percentile range across seeds. These are algorithmic sensitivity ranges, not confidence intervals.

Table A3. Fifty-seed sensitivity, $K = 4$ and $K = 6$ (gap = OTA − operator prevalence).

K	Topic (Top Terms, Seed 42)	Median Gap	5th–95th	Sign Stability
4	travel, companies, services, service, platforms, search, providers, online	+0.169	[−0.145, +0.375]	76%
4	new, technologies, capabilities, risks, technology, companies, financial, data	−0.006	[−0.352, +0.315]	50%
4	risks, business, services, generative, technologies, results operations, privacy	−0.052	[−0.380, +0.209]	64%
4	generative, travel, solutions, generative ai, use, marketing, artificial intelligence	−0.057	[−0.347, +0.231]	60%
6	travel, companies, services, search, service, platforms, providers, travel service	+0.116	[−0.164, +0.342]	72%
6	increase, financial, business financial, business, risks, operations, financial condition	+0.069	[−0.176, +0.239]	70%
6	use, uncertainty, rights, intellectual property, validity, enforceability	+0.051	[−0.108, +0.244]	66%
6	new, technologies, business, financial, risks, capabilities, adversely affect	−0.018	[−0.237, +0.264]	54%
6	risks, business, services, technologies, generative, privacy, regulatory	−0.063	[−0.369, +0.127]	68%
6	generative, travel, solutions, generative ai, marketing, use, generative artificial	−0.086	[−0.372, +0.075]	84%

Every 5th–95th range crosses zero, no topic keeps a consistent gap sign in more than 84 percent of seeds, and a firm-level permutation of the side labels yields $p = 0.114$ for the largest observed gap; leave-one-firm-out gaps for the travel-search topic range from +0.04 to +0.34. Recurring themes are recognizable across seeds, but the partition and the side prevalences are algorithmically unstable, which is why Section 4.2 rests on the manual thematic coding and the topic model is reported here as corroboration only. Corresponding 50-seed tables for $K = 2$, 3, and 5 accompany the replication files.

Appendix C. Convergence Series and Non-AI Baselines

Table A4. Full convergence series, AI-related language (raw).

Fiscal Year	Mean Pairwise Cosine	Firms Contributing
FY2019	0.189	4
FY2020	0.117	4
FY2021	0.063	4
FY2022	0.061	4
FY2023	0.081	6
FY2024	0.125	7
FY2025	0.212	7

Table A4 reports the full convergence series and Table A5 the single-estimand baseline comparison underlying Figure 1. Pre-2023 values rest on very few AI passages per firm (the anchored corpus is dominated by generic AI mentions before the GenAI period) and should be read as noisy; the GenAI-period rows are those interpreted in Section 4.3.

Table A5. Hotel–OTA differences in risk-sentence composition, FY2023–2025 (single estimand).

Category	Operator Share	OTA Share	$\lvert d \rvert$
GenAI (strict anchor rule)	0.67%	1.32%	0.51
Cybersecurity/data privacy	7.68%	5.24%	0.68
Macroeconomy/travel demand	5.21%	3.58%	0.97
Workforce/labor	11.40%	6.55%	1.88

Notes: Each cell is the mean share of a firm-year's risk-factor sentences matching the category dictionary, FY2023–2025, firm-years weighted equally; $\lvert d \rvert$ is the absolute standardized difference across firm-years. Non-AI baseline categories differ across the dyad by as much as or more than GenAI incidence does, which is why the paper's differentiation claim (Section 4.2) rests on the thematic composition of GenAI language rather than on GenAI being uniquely differentiated.

Appendix D. Codebook, Firm Panel, and Coding Reliability

Table A6 reports firm-level theme shares and Table A7 reports the theme-by-stance alignment. Thematic codebook. (i) *AI-mediated discovery and disintermediation*: GenAI relocating travel search, discovery, or booking away from the firm's channel—AI assistants or agents booking directly, platforms satisfying user intent without pass-through, LLM entry into travel booking, supplier-direct enablement. (ii) *Rival capability asymmetry*: competitors or large technology firms using GenAI to compete more effectively, absent the channel-relocation mechanism—innovation speed, resource asymmetry, replication of functionality, competitor enumerations. (iii) *Own deployment and adoption*: the firm deploys, incorporates, or invests in GenAI, including the talent and infrastructure demands of adoption. (iv) *Deployment liabilities and compliance*: legal, regulatory, intellectual-property, privacy, security, or ethical risks from GenAI use, including AI-enabled fraud and hacking. (v) *General/descriptive*: ambient uncertainty, factor lists, and information-availability statements without a specific channel or deployment claim.

Table A6. Firm-level theme shares.

Firm	Side	Discovery	Rival	Deployment	Liability	General	*N*
Booking Holdings	OTA	0.26	0.17	0.19	0.31	0.07	42
Expedia Group	OTA	0.00	0.67	0.17	0.17	0.00	6

Table A6. *Cont.*

Firm	Side	Discovery	Rival	Deployment	Liability	General	*N*
Tripadvisor	OTA	0.00	0.00	1.00	0.00	0.00	3
Hyatt Hotels	Operator	0.17	0.00	0.50	0.17	0.17	6
Hilton Worldwide Holdings	Operator	0.14	0.00	0.00	0.86	0.00	7
InterContinental Hotels Group	Operator	0.11	0.00	0.00	0.56	0.33	9
Marriott International	Operator	0.00	0.00	0.00	1.00	0.00	2

Table A7. Theme × stance alignment (modal labels).

Theme	Offensive	Threat	Neutral
AI-mediated discovery and disintermediation	0	11	3
Rival capability asymmetry	0	11	0
Own deployment and adoption	8	2	5
Deployment liabilities and compliance	0	20	8
General/descriptive	0	3	4

Stance coding reliability. Five passes (local qwen2.5:14b-instruct, temperature 0.2, seeds 101/211/307/401/503), each under an independently randomized passage order: mean pairwise Cohen's $\kappa = 0.68$ (minimum 0.62), mean agreement 82.9 percent (minimum 78.7), 50 of 75 passages unanimous, 11 at four-of-five, 14 at three-to-two splits, no unresolved ties. A one-passage-at-a-time pass (seed 601) agreed with the five-pass modal label on 90.7 percent of passages. The 14 split passages were re-examined and adjudicated against the coding rule; the adjudication log records each item's five votes, final label, and rationale, and Section 4.4 reports results under both label sets. The exact prompts and pass orders are provided in the Supplementary Materials; the raw model responses are available from the corresponding author on request.

Appendix E. Underpowered and Unidentified Causal Tests

The following supplementary analyses are reported for transparency. They are not treated as evidential, and the text cautions explicitly against reading them as evidence of absence.

Event studies. Cumulative abnormal returns for OTAs around major AI product launches were estimated under market-model, Fama–French three-factor and Carhart four-factor specifications. Group-level effects were statistically indistinguishable from permutation (placebo) distributions across all specifications. With three OTAs and short event windows, the design lacks the power to detect economically plausible effects; the estimates are uninformative rather than null.

Difference-in-differences. A two-way fixed-effects specification on OTA take-rate is not identified. The GenAI shock arrives contemporaneously for all OTAs in 2023, so the post-treatment indicator is collinear with year fixed effects and no untreated control group exists. This is a structural feature of a common shock rather than an empirical finding, and no coefficient from such a specification should be interpreted.

Synthetic control. A synthetic-control exercise for a single OTA around its own AI product launch produced a post-period divergence that was not distinguishable from placebo under in-space inference.

Taken together, these results reflect the structural limits of a small, common-shock sample rather than evidence that GenAI has had no economic effect on intermediation. They motivate the identification strategies proposed in Section 5.3—in particular, exploiting staggered, firm-specific AI events in order to recover a credible control group.

References

1. Toh, R.S.; Raven, P.; DeKay, F. Selling Rooms: Hotels vs. Third-Party Websites. *Cornell Hosp. Q.* **2011**, *52*, 181–189. [CrossRef]
2. Buhalis, D.; Law, R. Progress in Information Technology and Tourism Management: 20 Years On and 10 Years after the Internet—The State of eTourism Research. *Tour. Manag.* **2008**, *29*, 609–623. [CrossRef]
3. Law, R.; Leung, R.; Lo, A.; Leung, D.; Fong, L.H.N. Distribution Channel in Hospitality and Tourism: Revisiting Disintermediation from the Perspectives of Hotels and Travel Agencies. *Int. J. Contemp. Hosp. Manag.* **2015**, *27*, 431–452.
4. Anderson, C.K. The Billboard Effect: Online Travel Agent Impact on Non-OTA Reservation Volume. *Cornell Hosp. Rep.* **2009**, *9*, 6–9.
5. Anderson, C.K.; Han, S. The Billboard Effect: Still Alive and Well. *Cornell Hosp. Rep.* **2017**, *17*, 3–10.
6. O'Connor, P.; Assaker, G.; El Haddad, R. Online Travel Agency Participation: An Empirical Investigation of Its Financial Contribution to U.S. Hotel Profitability. *Cornell Hosp. Q.* **2025**, *66*, 527–538. [CrossRef]
7. Ivanov, S.; Webster, C. Robots in Tourism: A Research Agenda for Tourism Economics. *Tour. Econ.* **2020**, *26*, 1065–1085. [CrossRef]
8. Tussyadiah, I. A Review of Research Into Automation in Tourism: Launching the Annals of Tourism Research Curated Collection on Artificial Intelligence and Robotics in Tourism. *Ann. Tour. Res.* **2020**, *81*, 102883. [CrossRef]
9. Eisfeldt, A.L.; Schubert, G.; Zhang, M.B.; Taska, B. Generative AI and Firm Values. *J. Finance*, 2026; *in press*. [CrossRef]
10. Babina, T.; Fedyk, A.; He, A.; Hodson, J. Artificial Intelligence, Firm Growth, and Product Innovation. *J. Financ. Econ.* **2024**, *151*, 103745. [CrossRef]
11. Entman, R.M. Framing: Toward Clarification of a Fractured Paradigm. *J. Commun.* **1993**, *43*, 51–58. [CrossRef]
12. Dutton, J.E.; Jackson, S.E. Categorizing Strategic Issues: Links to Organizational Action. *Acad. Manag. Rev.* **1987**, *12*, 76–90. [CrossRef] [PubMed]
13. Gilbert, C.G. Unbundling the Structure of Inertia: Resource Versus Routine Rigidity. *Acad. Manag. J.* **2005**, *48*, 741–763. [CrossRef]
14. Bakos, J.Y. Reducing Buyer Search Costs: Implications for Electronic Marketplaces. *Manag. Sci.* **1997**, *43*, 1676–1692. [CrossRef]
15. Rochet, J.-C.; Tirole, J. Platform Competition in Two-Sided Markets. *J. Eur. Econ. Assoc.* **2003**, *1*, 990–1029. [CrossRef]
16. Campbell, J.L.; Chen, H.; Dhaliwal, D.S.; Lu, H.-M.; Steele, L.B. The Information Content of Mandatory Risk Factor Disclosures in Corporate Filings. *Rev. Account. Stud.* **2014**, *19*, 396–455. [CrossRef]
17. Kravet, T.; Muslu, V. Textual Risk Disclosures and Investors' Risk Perceptions. *Rev. Account. Stud.* **2013**, *18*, 1088–1122. [CrossRef]
18. Dyer, T.; Lang, M.; Stice-Lawrence, L. The Evolution of 10-K Textual Disclosure: Evidence from Latent Dirichlet Allocation. *J. Account. Econ.* **2017**, *64*, 221–245. [CrossRef]
19. Blei, D.M.; Ng, A.Y.; Jordan, M.I. Latent Dirichlet Allocation. *J. Mach. Learn. Res.* **2003**, *3*, 993–1022.
20. Brown, S.V.; Tucker, J.W. Large-Sample Evidence on Firms' Year-over-Year MD&A Modifications. *J. Account. Res.* **2011**, *49*, 309–346. [CrossRef]
21. Hassan, T.A.; Hollander, S.; van Lent, L.; Tahoun, A. Firm-Level Political Risk: Measurement and Effects. *Q. J. Econ.* **2019**, *134*, 2135–2202. [CrossRef]
22. Loughran, T.; McDonald, B. When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks. *J. Finance* **2011**, *66*, 35–65. [CrossRef]
23. Loughran, T.; McDonald, B. Measuring Readability in Financial Disclosures. *J. Finance* **2014**, *69*, 1643–1671. [CrossRef]
24. Loughran, T.; McDonald, B. Textual Analysis in Accounting and Finance: A Survey. *J. Account. Res.* **2016**, *54*, 1187–1230. [CrossRef]
25. Li, J.; Feng, Y.; Li, G.; Sun, X. Tourism Companies' Risk Exposures on Text Disclosure. *Ann. Tour. Res.* **2020**, *84*, 102986. [CrossRef] [PubMed]
26. Ling, L.; Guo, X.; Yang, C. Opening the Online Marketplace: An Examination of Hotel Pricing and Travel Agency On-Line Distribution of Rooms. *Tour. Manag.* **2014**, *45*, 234–243. [CrossRef]
27. Guo, X.; He, L. Tourism Supply-Chain Coordination: The Cooperation Between Tourism Hotel and Tour Operator. *Tour. Econ.* **2012**, *18*, 1361–1376. [CrossRef]
28. Nicolau, J.L.; Sharma, A. To Ban or Not to Ban Rate Parity, That Is the Question…or Not? *Int. J. Hosp. Manag.* **2019**, *77*, 523–527. [CrossRef]
29. Balyuk, T.; Davydenko, S. Reintermediation in FinTech: Evidence from Online Lending. *J. Financ. Quant. Anal.* **2024**, *59*, 1997–2037. [CrossRef]
30. Li, F. The Information Content of Forward-Looking Statements in Corporate Filings—A Naïve Bayesian Machine Learning Approach. *J. Account. Res.* **2010**, *48*, 1049–1102. [CrossRef]
31. Staw, B.M.; Sandelands, L.E.; Dutton, J.E. Threat-Rigidity Effects in Organizational Behavior: A Multilevel Analysis. *Adm. Sci. Q.* **1981**, *26*, 501–524. [CrossRef]
32. Gursoy, D.; Li, Y.; Song, H. ChatGPT and the Hospitality and Tourism Industry: An Overview of Current Trends and Future Research Directions. *J. Hosp. Mark. Manag.* **2023**, *32*, 579–592. [CrossRef]

33. Dwivedi, Y.K.; Pandey, N.; Currie, W.; Micu, A. Leveraging ChatGPT and Other Generative Artificial Intelligence (AI)-Based Applications in the Hospitality and Tourism Industry: Practices, Challenges and Research Agenda. *Int. J. Contemp. Hosp. Manag.* **2024**, *36*, 1–12. [CrossRef]
34. Dogru, T.; Line, N.; Mody, M.; Hanks, L.; Abbott, J.; Acikgoz, F.; Assaf, A.; Bakir, S.; Berbekova, A.; Bilgihan, A.; et al. Generative Artificial Intelligence in the Hospitality and Tourism Industry: Developing a Framework for Future Research. *J. Hosp. Tour. Res.* **2025**, *49*, 235–253. [CrossRef]
35. Wong, I.A.; Lian, Q.L.; Sun, D. Autonomous Travel Decision-Making: An Early Glimpse Into ChatGPT and Generative AI. *J. Hosp. Tour. Manag.* **2023**, *56*, 253–263. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.