



Article

How Organizations Choose Open-Source Generative AI Under Normative Uncertainty: The Moderating Role of Exploitative and Exploratory Behaviors

Suengjae Hong, Hakshun Ryee, Xiaoyan Jin  and Daegyul Yang * 

Department of Business Administration, School of Management, Kyung Hee University, Seoul 02447, Republic of Korea

* Correspondence: daegyul@khu.ac.kr; Tel.: +82-2-961-2233

Abstract

Open-source generative AI technologies offer transparent and customizable alternatives to proprietary AI systems, the concept of which closely aligns with the principles of open innovation. Organizations with strong open-source orientations may have greater absorptive capacity to adopt open-source generative AI technologies. However, adopting such technologies into the organizations is not always guaranteed because ethical, privacy, and regulatory concerns on open-source generative AI usage create normative uncertainty that can reduce organizations' willingness to adopt the technology, particularly when it is used in customer-facing products or services rather than integrated into internal processes. This study draws on organizational learning theory and open innovation literature to examine how open-source orientation affects open-source generative AI adoption under normative uncertainty, and how this relationship depends on organizational exploiting and exploring behaviors. Using global survey data from the Linux Foundation, we test our hypotheses with ordered logistic regression and interaction effects. The results show that open-source oriented organizations are more likely to adopt open-source generative AI, but this effect weakens when normative uncertainty is high, especially in product-related use cases. These findings extend absorptive capacity theory by highlighting ethical ambiguity as a key moderating factor and provide practical insights into how organizations can responsibly approach open-source generative AI adoption.

Keywords: open-source generative AI; normative uncertainty; organization theory; organizational learning; exploration; exploitation; absorptive capacity; open innovation



Academic Editor: Albert Y. S. Lam

Received: 30 June 2025

Revised: 15 August 2025

Accepted: 12 September 2025

Published: 16 September 2025

Citation: Hong, S.; Ryee, H.; Jin, X.; Yang, D. How Organizations Choose Open-Source Generative AI Under Normative Uncertainty: The Moderating Role of Exploitative and Exploratory Behaviors. *J. Theor. Appl. Electron. Commer. Res.* **2025**, *20*, 250. <https://doi.org/10.3390/jtaer20030250>

Copyright: © 2025 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 (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

While a growing body of research has explored the transformative potential of generative AI for enhancing innovation, efficiency, and strategic decision-making across organizational domains [1–3], relatively little attention has been given to the unique organizational dynamics surrounding the adoption of open-source generative AI. Moreover, little is known about how normative uncertainty manifests differently depending on the application context of open-source generative AI, namely, whether it is integrated into internal process optimization or deployed for external product and service innovation. Therefore, drawing on organizational learning theory and the literature on open innovation, this study focuses on open-source oriented organizations that are considering the adoption of open-source generative AI and examines how their responses to a normatively uncertain environment may moderate their intention to adopt such technologies. By engaging in developer communities and institutionalizing favorable internal governance practices around open-source,

open-source oriented organizations facilitate stronger learning routines and, consequently, greater absorptive capacity [4,5]. This enables them to better recognize, assimilate, and apply the external knowledge embedded in open-source generative AI models. In this sense, the adoption of open-source generative AI can be seen as an example of external, community-driven knowledge that is integrated into organizational boundaries [4]. Its value proposition aligns closely with the broader open innovation paradigm [6], in which knowledge flows across organizational boundaries, vendor lock-in is mitigated, adaptive experimentation is enhanced [7,8], so that innovation becomes increasingly community-driven, cumulative, and distributed through developer community [9,10]. This dynamic is enabled not only by greater absorptive capacity, but also by desorptive capacity, through which organizations externalize internal insights and participate in the co-evolution of open-source ecosystems. These mutually reinforcing capabilities help open-source oriented organizations internalize community-driven knowledge such as open-source generative AI into not only their internal operations, but also contributing back to its development community, thus creating a virtuous cycle of distributed innovation.

However, while open-source oriented environments foster such learning behaviors, their influence is not unconditional. open-source generative AI entails unvalidated risks, including not only technological concerns but also ethical ones, particularly those related to privacy, security, and regulatory compliance, many of which remain insufficiently understood or operationalized. These ethical concerns generate normative uncertainty surrounding the adoption of open-source generative AI within organizations. Prior research on open innovation under such uncertainty suggests that ambiguity can compromise the legitimacy of absorbing external knowledge, even in organizations that are generally supportive of using open-source software [7]. We argue that the effect of normative uncertainty is contingent on the application context of open-source generative AI, especially, whether open-source generative AI models are embedded into internal processes or integrated into external-facing products. When open-source generative AI is deployed for internal use cases (e.g., workflow automation, content or code generation, internal analytics), organizations retain stronger control over data use and system integration. This setting reflects an exploitative learning logic, where open-source generative AI is used to refine and optimize existing routines [11], thereby weakening the perceived salience of normative uncertainty. In contrast, when open-source generative AI is deployed for customer-facing products or services, organizations are more vulnerable to reputational damage and regulatory intervention if any misuse occurs [12,13]. These exploratory contexts involve greater environmental complexity and higher stakes for legitimacy, amplifying the negative effect of normative uncertainty on open-source generative AI adoption intentions [14,15].

The empirical context of this study is based on a global survey conducted by the Linux Foundation, which investigates organizational attitudes and adoption behaviors toward open-source generative AI. We examine how organizational orientation toward open-source affects the likelihood of open-source generative AI adoption and how this relationship is shaped by the perceived ethical and regulatory uncertainties surrounding these technologies. Drawing on organizational learning theory and the open innovation framework, we argue that open-source oriented organizations tend to have greater absorptive capacity to engage with community-driven technologies like open-source generative AI. However, we propose that this positive orientation may be attenuated under conditions of normative uncertainty, particularly, when concerns about privacy, security, and compliance are salient. Furthermore, we suggest that the strength of this attenuated effect is contingent upon the application context of open-source generative AI. Specifically, the restraining effect of normative uncertainty on open-source generative AI adoption is expected to be weaker when organizations tend to adopt Gen AI models into internal

processes and stronger when they are deployed in external-facing products or services. To operationalize normative uncertainty, this study employs a supervised learning approach combining hierarchical clustering and principal component analysis on survey responses related to ethical, privacy, and regulatory concerns. This method allows us to capture the latent construct of normative uncertainty more precisely than simple summation of indicators, enabling nuanced empirical testing of its moderating role. We empirically tested our hypotheses using ordered logistic regression with interaction terms, showing how organizations manage the risks and opportunities of adopting open-source generative AI, depending on usage situation.

This study contributes to the literature on open innovation and organizational learning by articulating how normative uncertainty moderates the adoption of open-source generative AI within open-source oriented organizations, and how such effects vary depending on the application context, whether open-source generative AI is deployed for internal exploitation or external exploration. The findings offer both theoretical and practical implications: theoretically, we extend absorptive capacity frameworks to normatively contested technologies, and practically, we inform technology leaders and policymakers on how to mitigate barriers to responsible AI adoption. The remainder of this article is structured as follows. Following the introduction, the next part reviews the relevant theoretical foundations, including organizational learning and open-source innovation. This is followed by the development of the research hypotheses and conceptual model. The subsequent part outlines the research design and data collection process. Empirical results are then presented, followed by a discussion of the findings. The article then elaborates on the theoretical and managerial implications of the study and concludes with limitations and suggestions for future research.

2. Hypothesis Development

2.1. Literature Review

2.1.1. Absorptive Capacity and Desorptive Capacity in Open Innovation

In an open innovation context, in which organizations purposely use inflow and outflow of knowledge, absorptive capacity underpins inbound innovation (outside-in knowledge flows) and desorptive capacity underpins outbound innovation (inside-out knowledge flows). Absorptive capacity refers to an organization's ability to recognize the value of new external knowledge and plays a crucial role by allowing organizations to internalize external ideas and technologies [4,5,16]. A high absorptive capacity enhances an organization's ability to scan technological and market environments, identify valuable R&D partnerships, and integrate external innovations such as open-source technologies. Prior research emphasizes that absorptive capacity is strengthened through sustained R&D investment and organizational learning routines, the capacity of which is critical for leveraging inbound open innovation opportunities, as they enable organizations to recognize, assimilate, and apply external knowledge more effectively [4,5,17]. Unlike absorptive capacity, desorptive capacity is the complementary, outward-oriented capability that enables organizations to identify opportunities for external technology transfer and to facilitate the effective application of that knowledge at external recipients [18]. It plays a crucial role in outbound open innovation strategies by allowing organizations to unlock the value of internally developed or retained knowledge through external diffusion channels [19,20]. Empirical studies show that organizations with high desorptive capacity are more likely to engage in successful out-licensing, strategic partnerships, and open-source collaboration, as they are better equipped to match internal knowledge with external demand and ensure its usability [19,21,22]. In other words, desorptive capacity reflects an organizational ability to externalize or export knowledge by sharing innovations with R&D partners, licensing pro-

proprietary technologies, or contributing to open-source projects, thus reinforcing its position in broader innovation ecosystems.

Combining these two concepts, absorptive capacity and desorptive capacity can be seen as co-evolving theoretical constructs in open innovation management, both of which are essential for organizations operating within the open-source ecosystem [6,20,23]. From this perspective, this study argues that organizations with a stronger orientation toward open-source (open-source oriented organizations) are more likely to have established the absorptive and desorptive capacities necessary for pursuing new innovations related to open-source. Given that open innovation requires value “creation and exchange” [20] through interorganizational learning processes [22,24], open-source oriented organizations are more likely to develop comparable knowledge structures or to engage in collaborative interactions with other organizations that share similar open innovation value [25]. Consequently, open-source oriented organizations are better positioned to make strategic choices in favor of open-source innovations, thereby increasing their propensity to embrace a new open innovation.

2.1.2. Normative Uncertainty and Intentions to Adopt New Open Innovations

This study posits that open-source oriented organizations may be discouraged from adopting open-source innovation if they perceive it to be normatively uncertain. Given that normative uncertainty involves ambiguity over whether a given innovation conforms to prevailing societal norms, values, and expectations, it reflects the absence of stakeholder consensus about the legitimacy of an innovation or a new technology practice within a social or cultural context [26,27]. In particular, normative uncertainty prevails in the early stages of emerging technologies or institutional fields in which evaluative standards are still fluid and contested [15,28–30].

Under conditions of market uncertainty, absorptive capacity shapes the formation of expectations, enabling organizations to more accurately anticipate the nature and commercial potential of technological advances [4]. However, ambiguity stemming from normative uncertainty surrounding a new form of open innovation technology may increase an organization’s risk aversion toward adopting the technology because interpreting ambiguous signals is a central challenge in adaptive routines [31], and when normative legitimacy is unclear, this challenge intensifies, often leading to delayed or suboptimal adaptive responses [32] and shaping the contextual conditions under which innovations are adopted and diffused within organizations [28].

Normative uncertainty influences not only an organization’s absorptive capacity, but also its desorptive capacity, which involves both identifying external knowledge exploitation opportunities and transferring such knowledge to external recipients, guided by the organization’s monetary and strategic motives [33]. With its ambiguous signals, normative uncertainty can obscure the perceived value or appropriateness of knowledge transfer, complicate the identification of exploitable opportunities, and undermine the organization’s ability to leverage prior knowledge effectively [34,35]. As a result, organizations may be less able to externalize knowledge efficiently when normative legitimacy is unclear. Moreover, organizations may face a greater likelihood of legitimacy loss if stakeholders ultimately reject the adopted innovation as socially or ethically misaligned [36]. Empirical research in organizational theory shows that ethical or normative controversies often trigger more severe and enduring performance declines than purely economic setbacks, due to reputational damage, disrupted stakeholder relationships, and increased regulatory scrutiny [13,29,37–39].

2.1.3. The Impact of Different Application Contexts on New Open Innovation Adoption

While normative uncertainty can act as a mitigating factor in adoption decisions, its effect is not uniform across technology application contexts. From an exploration and exploitation perspective [11], internal process improvements generally correspond to exploitation logic. These improvements target already established processes with relatively clear goals, allowing organizations to refine, reuse, and enhance existing knowledge and routines while managing risks within well-defined boundaries [40,41]. Even if normative uncertainty generates negative impacts on the intention to adopt a new open innovation, its impact is largely contained within the organization or the scope of the tasks to which the technology is applied.

In contrast, externally facing products and services reflect an exploration logic, involving novel applications, experimentation, and engagement with external stakeholders [42]. The implications of normative uncertainty extend beyond internal performance. Negative signals can propagate to external stakeholders, potentially affecting reputation, customer trust, and broader social legitimacy. The risks associated with normative uncertainty in these contexts are therefore both greater in scope and more complex, requiring organizations to manage ambiguous outcomes under higher scrutiny and accountability pressures [14]. Accordingly, while open innovation adoption in internal process improvement contexts may be only modestly constrained by normative uncertainty, adoption for external-facing products or services is expected to be much more sensitive to such uncertainty.

2.2. Hypotheses

2.2.1. Intention of Open GenAI Adoption by Open-Source Oriented Organizations

As generative AI technologies evolve rapidly, organizations increasingly confront strategic decisions about how to adopt these technologies into their system. generative AI refers to artificial intelligence systems that can produce content, including human-like text, images, code, audio, etc., by learning generative patterns of content from vast data. Technically, generative AI represents a leap in AI capability, as modern large-language models (e.g., GPT-series by Open AI, LLaMA-series by Meta) are built on advanced neural architectures and trained on trillions of words from a vast of documents, so that those models are able to generate contextually relevant output. Unlike previous generations of AI, which primarily emphasized computational efficiency and automation, the value of generative AI lies in its ability to augment creativity, cognition, and human-like reasoning. For example, generative AI can assist in writing code, drafting documents, designing graphics, or even discovering molecules [43]. As the markets for generative AI tools grow exponentially, hundreds of startups and many enterprises have integrated generative AI into their products and workflow [1]. However, the development and deployment of generative AI have traditionally been dominated by proprietary systems because building high-performing generative AI models requires vast computational resources, access to massive datasets, and highly specialized expertise, conditions of which only a few well-funded tech companies could satisfy. For that reason, until recently, generative AI has remained a centralized and capital-intensive technology, governed by a small number of organizations with the infrastructure to train and operate such models. This proprietary concentration has led to growing concerns about transparency, accountability, and innovation bottlenecks, particularly as generative AI becomes increasingly integrated into core organizational workflows. In response to these limitations, a parallel movement has emerged. open-source generative AI, a new frontier of AI movements, promises to democratize access to cutting-edge generative models and fosters a more collaborative, transparent, and inclusive innovation ecosystem.

The emergence of open-source generative AI has marked a paradigmatic shift in how organizations adopt, develop, and appropriate artificial intelligence technologies. Unlike proprietary AI, open-source generative AI projects enable organizations to access, fine-tune, and deploy models with greater autonomy, transparency, and cost-efficiency. Not just such technical accessibility, open-source generative AI allows organizations to build upon open-sourced models that can be developed collaboratively through open-source communities, the result of which fosters innovation and technological improvement. This openness brings several key advantages for organizations considering the adoption of open-source generative AI. First, open-source generative AI significantly enhances accessibility and cost efficiency in model development and adoption in organization. By providing openly available models and weights, open-source generative AI dramatically lowers the barriers for researchers, startups, and organizations to access and utilize generative AI models without requiring big-tech-level infrastructure. For instance, open-source generative AI models can be deployed and fine-tuned locally, so that users can avoid the high costs and restrictions associated with proprietary APIs offered by major technology organizations. Second, open-source generative AI enables a collaborative innovation ecosystem, as the open-source nature allows developer communities to jointly contribute to model development and refinement by sharing data, techniques, and best practices. This collaborative approach facilitates the sharing of knowledge and of best practices, so that spreads the R&D costs and accelerates the developmental process. Moreover, aligning with the ethos of community-driven innovation, multiple stakeholders, including researchers, startups, and many organizations, can co-create the future of AI in an inclusive, transparent, and equitable manner. Third, open-source generative AI fosters transparency, governance, and trust. Because both source code and model weights are inspectable, open-source generative AI can be audited for biases or flaws, safety issues, and accountability by the open-source communities. For instance, the movement toward responsible AI, which can be defined as “an approach to developing, assessing, and deploying AI systems in a safe, trustworthy, and ethical way” [44], offers a governance framework that can flexibly address ethical, regulatory, and safety concerns on open-source generative AI adoption [45]. Furthermore, open-source generative AI mitigates vendor lock-in and promotes interoperability between established systems and AI models, so that organizations easily integrate open-source generative AI technologies across diverse systems and platforms. In sum, the open-source approach encourages organizations to adopt and to experiment with generative AI by making it more accessible, adaptable, and aligned with collaborative innovation practices.

Therefore, it is hypothesized that open-source oriented organizations, which have an open culture of using and contributing to open-source projects, are more likely to have higher absorptive and desorptive capacities to open-source generative AI technology. Organizations not only build absorptive capacity from the open-source community when they use external open-source software or technologies but also exercise desorptive capacity by contributing code or improvement back to the community via outward knowledge transfer. In an open-source generative AI context, open-source oriented organizations that have a high absorptive capacity can easily scan external knowledge, such as finding the latest open-source AI libraries or models and their structures and then, can internalize the knowledge. Such an organization with high absorptive capacity quickly recognizes the business value of a new open-source generative AI model and assimilates the knowledge into its R&D or products. Similarly, open-source oriented organizations tend to have processes for knowledge sharing. For example, they may contribute model fine-tuning techniques, bug fixes, or governance guidelines back to the open-source community, thereby exercising desorptive capacity through outward knowledge transfer. Open innovation literature indicates that organizations proficient in absorptive capacity and desorptive capacity gain more from

external collaborations, so that they are more successful in innovation performance [18]. Therefore, open-source oriented organizations likely tend to promote external knowledge scanning, to encourage employee participation in knowledge sharing, and to foster greater openness to integrating new open-source projects into their systems, all of which lower resistance to adopting external innovations, such as open-source generative AI, as examined in this study [9,46]. Therefore, we suggest the following hypothesis:

Hypothesis 1. *Organizations with a strong orientation toward open-source are more likely to intend to adopt open-source generative AI.*

2.2.2. Normative Uncertainty in Open-Source Generative AI Adoption

Normative uncertainty arises from unresolved ethical and societal concerns, such as data privacy, algorithmic fairness, and the potential misuse of generative content, if an organization considers adopting AI within its boundary. These concerns are particularly salient in the case of open-source generative AI, where the open-source nature of the technology limits centralized control and governance. Unlike proprietary systems, open-source generative AI allows broad access and modification by all access users, so that sometimes it may be difficult to ensure whether open-source generative AI models are consistent with ethical safeguards or compliance standards. For managers and policymakers, such openness creates both opportunities for innovation and challenges in establishing responsible deployment practices, so that it intensifies normative uncertainty surrounding open-source generative AI adoption in organizational settings.

Therefore, we argue that normative uncertainty can exert a mitigating effect on open-source oriented organizations' willingness to adopt open-source generative AI models. Although normative uncertainty, such as concerns about bias, hallucination, copyright infringement, privacy leakage, and workforce displacement have been raised in relation to proprietary generative AI [47], these concerns are further amplified in the case of open-source generative AI. For example, it raises questions about intellectual property and responsible use that legal ambiguities remain regarding licensing terms and the ownership of training data. Moreover, the open release of model weights increases the risk of misuse, as these models can be repurposed for harmful applications. Regulatory frameworks, such as the EU AI Act, have yet to establish clear guidelines for the governance of open-weight models, so that organizations are uncertain about compliance and accountability expectations [48]. While open-source oriented organizations often possess strong absorptive and absorptive capacities that enable them to scan, absorb, and disseminate external knowledge [4,18], these learning routines can still be affected by the uncertainty and confusion surrounding open-source generative AI. Organizational learning scholars argue that interpreting ambiguous signals is a central challenge in adaptive routines, and when normative legitimacy is unclear, such interpretive process becomes strained [31]. Research has shown that in contexts where normative uncertainty prevails, organizations tend to be more conservative and have risk-averse behaviors even if they have a history of open innovation [28]. Specifically, when organizations face unclear ethical, regulatory, or cultural norms around open-source generative AI use, such as licensing ambiguity, unclear data provenance, or fears of misuse, these doubts can inhibit their ability to absorb new knowledge and to share or contribute outwardly. Consequently, normatively uncertain contexts can reduce confidence in deploying learning routines that support open innovation. Therefore, we suggest the following hypothesis:

Hypothesis 2. *The likelihood that an open-source oriented organization adopts open-source generative AI is reduced when normative uncertainty is high.*

2.2.3. Process vs. Product Applications of Generative AI

When open-source generative AI models are adopted for internal use, such as automating workflows, enhancing document summarization, or supporting internal analytics, they are typically deployed within organizational systems [1,49,50]. These applications tend to be private to employees, easily reversible after deployment, and auditable through internal processes. Therefore, normative uncertainty is mainly associated with potential misuse or biased outputs, and its risk is limited within the organizational boundary. Therefore, normative uncertainty associated with the use of open-source generative AI is often managed through internal controls and is less likely to engage reputational risks or regulatory scrutiny. Consistent with exploitative learning, organizations can incrementally adapt their routines to open-source generative AI in these internal use cases. By contrast, deploying open-source generative AI in products or services, such as customer-facing chatbots, marketing content generators, or personalized recommendation engines, can expose organizations to substantially greater reputational, legal, and ethical risks. In such exploratory contexts, negative feedback from markets or regulators may be more diffuse, accountability is often externally imposed, and the consequences of errors are typically amplified [51]. Therefore, under normative uncertainty, organizations should consider not only ensuring functional accuracy but also justifying the ethical soundness of gen AI behaviors to customers, regulators, and other audiences in markets. Consequently, the negative moderating effect of normative uncertainty on the intention of open-source oriented organizations to adopt open-source generative AI is conditional upon the application contexts. In short, in exploitation contexts, the impact of normative uncertainty is buffered by bounded organizational controls and limited visibility to external stakeholders, while in exploration contexts, the dampening effect of normative uncertainty on the intention to adopt generative AI can be amplified by the unpredictability of outcomes for markets and regulators. Therefore, we propose the following hypotheses, which together form a three-way interaction that contextualizes the organizational learning dynamics underlying the adoption of open-source generative AI.

Hypothesis 3a. *The negative moderating effect of normative uncertainty on the relationship between open-source orientation and open-source generative AI adoption is attenuated when open-source generative AI is applied to internal organizational processes.*

Hypothesis 3b. *The negative moderating effect of normative uncertainty on the relationship between open-source orientation and open-source generative AI adoption is amplified when open-source generative AI is applied to external products or services.*

3. Methods

3.1. Data

To empirically examine how organizational orientation toward open-source influences the intention of Open GenAI adoption under conditions of normative uncertainty, we used a cross-sectional survey dataset. The data were drawn from a global survey conducted by the Linux Foundation in 2024, which gathered responses from industry-specific companies, IT vendors and service providers, as well as nonprofit, academic, and government organizations, spanning the Americas, Europe, Asia-Pacific, and other regions, and representing firms from startups to large enterprises. Respondents reported on a diverse set of organizational use cases for GenAI, covering applications in technology, manufacturing, finance, healthcare, public services, and other domains. The survey collected data on organizational characteristics (industry, size, location), current GenAI adoption and

investment levels, ethical, legal, and governance concerns, the extent of open-source AI technology and infrastructure use, specific GenAI techniques applied in organizations, and the benefits organizations have realized from GenAI. Recruitment drew from Linux Foundation subscribers, members, partner communities, and social media channels. To ensure data validity, the survey incorporated multiple prescreening and screening stages as follows: respondents were first qualified based on their professional experience and their familiarity with GenAI adoption within their organization, and additional screening questions and data quality checks were applied to remove ineligible or low-quality responses. These procedures ensured that final respondents could reliably answer on behalf of their organization. Eligible respondents were required to be familiar, very familiar, or extremely familiar with their organization's GenAI adoption strategy and status and to hold senior professional roles, such as CTO, CIO, head of data science, or head of innovation/product management. A total of 316 respondents completed the survey before applying the eligibility and quality control criteria. To maintain analytical consistency and interpretability of ordinal outcome variables, we excluded cases where respondents selected non-substantive options such as "Don't know," "Not applicable," or provided incomplete responses on key variables. After applying these quality control and eligibility criteria, the final analytic sample consisted of 209 organizations.

3.2. Measurement

3.2.1. Dependent Variable: Intention to Adopt Open-Source Generative AI in Organization

To measure the dependent variable, we used the survey item: *"How do you expect this use of open-source generative AI in your organization to change in the next two years? (Select one)"* Respondents could select from six categories: (1) Don't know or not sure, (2) Substantially increase, (3) Increase, (4) Stay the same, (5) Decrease, and (6) Substantially decrease. For analytical clarity, responses of (1) Don't know or not sure were excluded from the analysis due to their ambiguous interpretability. The remaining response options reflect a directional and ordinal structure, ranging from strong positive expectations to strong negative expectations regarding future adoption. For ease of interpretation in our statistical models, we reverse-coded the responses such that higher values indicate stronger anticipated adoption of open-source generative AI technologies.

3.2.2. Independent Variables: Open-Source Orientation

To operationalize the independent variable for testing our first hypothesis, we relied on the survey item: *"How does the open-source nature of a tool or model influence its adoption within your organization? (Select one)"* Participants selected one response from six ordered categories: (1) Don't know or not sure, (2) Strongly negative: We avoid open-source generative AI tools and models, (3) Negative: We are cautious about adopting open-source generative AI tools and models, (4) Neutral: The open-source nature does not affect our decision, (5) Positive: Being open-source is a favorable factor, and (6) Strongly positive: We prioritize open-source generative AI tools and models. To ensure interpretive clarity, responses marked as (1) were excluded from the analysis. The remaining responses reflect a graded scale of organizational openness toward open-source generative AI from strong skepticism to active preference. In our coding scheme, higher values represent a more favorable orientation toward open-source generative AI. To facilitate interpretation and reduce multicollinearity in interaction terms, the resulting Open-source Orientation variable was mean-centered prior to modeling.

3.2.3. Two-Way Moderating Variables: Normative Uncertainty

To operationalize the construct of normative uncertainty, we utilized two survey items that capture organizations' evaluative criteria and concerns in selecting and adopt-

ing generative AI tools, particularly those related to ethical ambiguity, privacy, security, and regulatory risk. These dimensions are theoretically grounded in prior literature on legitimacy tensions and contested norms in technological adoption [15].

The first item asked: “What are the most important characteristics your organization considers when choosing a generative AI model or tool? (Select all that apply)” This item contained 16 selectable attributes, including (1) Don’t know or not sure, (2) Does not apply to us, (3) Accuracy or performance, (4) Being open-source, (5) Compliance with regulations, (6) Cost, (7) Customizability, (8) Ease of integration, (9) Performance or speed, (10) Privacy, (11) Scalability, (12) Security, (13) Support and maintenance, (14) User experience, (15) Vendor reputation, (16) Other (please specify). The second item asked: “When adopting generative AI models and tools, what are your primary concerns? (Select all that apply).” This item comprised 25 selectable options, including: (1) Does not apply to us, (2) Cost of development, (3) Cost of operations, (4) Customization of the tools to meet our needs, (5) Deployment of the solution, (6) Ease of deployment, (7) Ease of use, (8) Ethical issues, (9) Model fine-tuning, (10) Integration with existing systems, (11) Lack of business needs, (12) Lack of skills or expertise, (13) Lack of support, (14) Latency of the models, (15) Privacy of our data, (16) Quality of AI output, (17) Regulatory compliance and legal uncertainties or liabilities, (18) Safety, (19) Security risks, (20) Technical challenges, (21) Technology maturity, (22) Trustworthy data and models, (23) Uncertain ROI, (24) Don’t know or not sure, and (25) Other (please specify).

To construct a latent measure of normative uncertainty, we conducted a two-step unsupervised learning procedure using agglomerative hierarchical clustering and principal component analysis (PCA). The analysis was implemented in Python 3.12 using the *scipy* and *scikit-learn* libraries. First, we prepared two sets of multi-response variables from the survey: one regarding the characteristics considered when choosing generative AI tools (13 binary items), and the other on primary concerns when adopting generative AI technologies (22 binary items). For each question, responses were dummy-coded into binary variables (1 = selected, 0 = not selected), and “Don’t know”, “Does not apply to us”, and “Other (specify)” responses were excluded for interpretive clarity, the latter of which contained non-standardized, open-text responses.

Next, we applied agglomerative hierarchical clustering to each question block to identify groups of co-occurring concerns. Specifically, we used the *scipy.cluster.hierarchy.linkage* function with the Ward method (minimizing intra-cluster variance) and calculated pairwise Euclidean distances between variables. To determine the optimal number of clusters (k), we used an elbow detection technique based on the second-order difference in linkage distances, implemented through a custom function using the linkage matrix. As a result of the hierarchical clustering analysis, we identified ten distinct clusters for each of the two multi-response survey modules: one is about the characteristics considered when selecting generative AI tools and the other addresses concerns about adopting generative AI models. Each cluster represents a set of survey items that exhibit similar response patterns across survey participants, and this indicates that items that tend to co-occur in response selections were grouped together, so that the survey items within a cluster reflects conceptual proximity or underlying latent dimensions among respondents. For that reason, this method is useful for constructing higher-order constructs, such as normative uncertainty, from a large set of binary indicators by capturing similarity in selection patterns empirically, not arbitrarily. Based on the results of the hierarchical clustering analysis, individual items were grouped into clusters as summarized in Table 1 (*Clustering of characteristics considered in Gen AI adoption*) and Table 2 (*Clustering of concerns considered in Gen AI adoption*).

Table 1. Clustering of Characteristics considered in Gen AI Adoption.

Cluster ID	Item Labels
Cluster 1	Ease of integration
Cluster 2	User experience
Cluster 3	Customizability, Support and maintenance, Vendor reputation
Cluster 4	Performance or speed
Cluster 5	Scalability
Cluster 6	Being open-source
Cluster 7	Accuracy or performance
Cluster 8	Cost
Cluster 9	Privacy, Security
Cluster 10	Compliance with regulations

Table 2. Clustering of Concerns considered in Gen AI Adoption.

Cluster ID	Item Labels
Cluster 1	Safety, Security risks
Cluster 2	Ethical issues, Privacy of our data, Regulatory compliance and legal uncertainties or liabilities
Cluster 3	Customization of the tools to meet our needs, Integration with existing systems
Cluster 4	Model fine-tuning
Cluster 5	Trustworthy data and models
Cluster 6	Quality of AI output
Cluster 7	Cost of development, Cost of operations
Cluster 8	Deployment of the solution, Ease of deployment, Ease of use
Cluster 9	Technical challenges, Technology maturity
Cluster 10	Lack of business needs, Lack of skills or expertise, Lack of support, Latency of the models, Uncertain ROI

Within the characteristics module, Cluster 9 and Cluster 10 included items related to Privacy, Security, and Compliance with regulations. Similarly, within the concerns module, Cluster 1 and Cluster 2 grouped together items associated with Safety, Security risks, Ethical issues, Privacy of our data, Regulatory compliance and legal uncertainties of liabilities. While clustering yielded a variety of interpretable theoretical constructs, our operationalization of normative uncertainty focused specifically on the clusters that reflect institutional, ethical, and socio-technical concerns, consistent with prior theoretical frameworks in organizational and innovation studies [27,52,53]. Accordingly, we selected and aggregated the aforementioned clusters (*Characters: Cluster 9 and Cluster 10; Concerns: Cluster 1 and Cluster 2*) to represent the latent construct of normative uncertainty, which was subsequently used in our models as a key explanatory factor. This construct was either aggregated through summation or further validated using dimensionality reduction techniques such as principal component or factor analysis, as detailed in the methodological Section 3.2. However, relying solely on summation has key limitations. Summing binary indicators implicitly assume equal weight and independent contribution of each item, so this approach ignores underlying correlations or shared variance among them. This may result in oversimplified or biased representations of latent constructs, particularly when multiple items cluster around the same conceptual dimension (e.g., privacy and security). In contrast, PCA or factor analysis accounts for these interdependencies and allows for more parsimonious and empirically grounded measurement.

Therefore, we aggregated the binary variables within each cluster to create composite variables representing clustered concern types. These cluster-based composites were then subjected to principal component analysis (PCA) using *sklearn.decomposition.PCA*. From

each block, the first principal component (PC1), which explained the highest proportion of variance, was extracted and standardized. The resulting principal component scores served as the operationalized measure of normative uncertainty in subsequent models. By transforming discrete dummy variables into a continuous latent dimension, this approach enables us to assess the extent to which an organization is sensitized to normative uncertainty. In other words, organizations that selected more items associated with ethical, privacy, and compliance-related concerns are likely to load higher on the principal component, thereby exhibiting greater perceived salience of normative risks surrounding open-source generative AI adoption, while reducing dimensionality and multicollinearity in the regression analysis. This continuous score captures variation in concern intensity across organizations and offers a refined measure for empirical modeling. (While factor analysis is a commonly employed method to extract latent constructs, its application assumes continuous, normally distributed data and linear relationships among variables. However, in our dataset, normative uncertainty was measured through multiple binary (dummy-coded) items derived from multi-response survey questions. Applying conventional factor analysis directly to such binary variables would violate these foundational assumptions, leading to biased factor loadings and misestimated communalities [54]. To address this, we adopted an alternative approach based on tetrachoric correlation, which is specifically designed to estimate the latent correlation between two dichotomous variables that are presumed to stem from underlying continuous distributions. We constructed a full tetrachoric correlation matrix for the binary items using a pairwise maximum-likelihood estimation method. Then, this matrix was analyzed with *FactorAnalyzer* to extract latent dimensions of normative uncertainty. Despite this adjustment, tetrachoric-based factor analysis also presents challenges. It is computationally intensive, particularly for large variable sets, and remains sensitive to sparse or imbalanced response patterns, which can lead to convergence issues or unstable estimates [55]). To ensure interpretability and mitigate multicollinearity in interaction models, normative uncertainty score was mean-centered before inclusion in regression analyses.

3.2.4. Three-Way Moderating Variables: Process- and Product-Oriented Application of Generative AI

To operationalize moderate variables for testing third and fourth hypotheses, we relied on the survey item: “How is your primary generative AI use case integrated into your business? (select one)” Participants selected one response from five categories: (1) Don’t know or not sure, (2) Too soon to tell, (3) Generative AI supports our internal processes, workflows, and tasks, (4) Generative AI is integrated into our products or services, and (5) We creating solutions that enable third parties to utilize generative AI in their products. This item was designed as a single-choice question, meaning each respondent could indicate only one dominant application context. As such, the responses reflect mutually exclusive commitments to either an exploitative or exploratory application of generative AI, consistent with March’s (1991) framework distinguishing between exploitation (internal optimization) and exploration (external innovation) [11].

We coded these responses to derive two binary moderator variables. Responses selecting category (3) were coded as 1 for the Process-Oriented (*exploitation*) variable and 0 otherwise. Similarly, responses selecting category (4) were coded as 1 for the Product-Oriented (*exploration*) variable and 0 otherwise. Other response options were excluded from the moderation analysis due to their interpretive ambiguity. This approach enabled us to examine how the effect of normative uncertainty on adoption intentions differs depending on whether open-source generative AI is primarily applied to internal processes or external-facing products.

3.2.5. Control Variables

Size. To construct the organizational size variable used as a control in the regression models, we relied on the survey item: “Approximately how many employees does your organization have worldwide? Your best estimate is fine.” Respondents selected a single category from the following options: (1) Don’t know or not sure, (2) Less than 10, (3) 10 to 249, (4) 250 to 999, (5) 1000 to 9999, (6) 10,000 to 19,999, (7) 20,000 or more. Responses in category (1) were excluded from the analysis due to their lack of interpretive value. The remaining responses were treated as a quasi-ordinal variable indicating increasing organizational size. For ease of interpretation and to retain the ordinal nature of the data, the variable was recorded from 1 to 6, where higher values correspond to larger organizations. This variable was included as a control to account for organizational capacity and structural complexity, which are known to influence both new technology adoption intentions and absorptive capabilities [5].

IT industry. To control for industry-level effects, we used the survey item: “Which option best describes your organization?” Respondents were asked to select one of five categories: (1) Don’t know or not sure, (2) I have not recently worked for an organization, (3) My organization primarily operates in the Information Technology (IT) sector, (4) My organization primarily operates outside the IT sector (e.g., education, finance, healthcare, retail, government, etc.), and (5) Other (please specify). Based on this item, we created a binary indicator variable for IT sector affiliation, coded as 1 if the respondent selected option (3), indicating the organization operates in the IT sector, and 0 otherwise. This control variable allows us to account for differences in adoption patterns that may arise from the sectoral context of the organization.

HQ Location. To account for regional effects, we constructed a binary variable indicating whether the organization is headquartered in North America. Based on the survey item “In which region does your organization have its primary headquarters?”, respondents selected a single category from the following options: (1) United States or Canada, (2) South America, Mexico, Central America, or the Caribbean, (3) Europe, (4) Asia-Pacific (except China, India, and Japan), (5) China, (6) India, (7) Japan, (8) Africa, (9) Middle East, (10) Other (please specify). We coded 1 if responses selected (1) United States or Canada, while all other regions were coded as 0. This allowed us to control for potential geographic differences in technology adoption environments, regulatory pressures, and open-source ecosystem maturity.

Open-source neutral hosting importance. To control for organizations’ general valuation of open-source infrastructure—beyond their orientation toward open-source tools themselves—we included an additional variable based on the item: “How important is it for your organization to use AI open-source tools that are hosted by a neutral party, such as the Linux Foundation, instead of a corporate entity?” Response options ranged from (1) Don’t know or not sure, (2) Not at all important, (3) Not important, (4) Somewhat important, (5) Important, (6) Very important, and (7) Extremely important. We excluded responses of Don’t know or not sure due to interpretive ambiguity.

This variable captures the strategic importance of neutral-hosted open-source tools, which may independently shape adoption decisions—for instance, organizations concerned with vendor lock-in or seeking greater governance transparency may emphasize the neutrality of hosting entities regardless of their internal open-source orientation [56]. However, since this variable is empirically correlated with open-source orientation, we addressed potential multicollinearity by orthogonalizing it. Specifically, we regressed this variable on the open-source orientation variable and retained the residuals as the final input. This process isolates the unique variance in the importance of neutral-hosted open-source tools

that are not explained by general open-source affinity, allowing us to better identify its independent effect on adoption outcomes.

Gen AI maturity stage. To account for the level of generative AI integration within organizations, we included a variable based on the item: “What is the current stage of your organization’s primary generative AI use case?” Respondents selected one of the following categories: (1) Don’t know or not sure, (2) Initial development: initial design and early testing, (3) Experimental deployment: evaluation in a controlled environment, (4) Stalled post-pilot: progress halted following initial testing phases, (5) Initial production deployment: launched for actual use within limited operational areas, and (6) Deployed at scale: fully implemented and operational. For analytical clarity, responses of “Don’t know or not sure” were excluded. The remaining categories were treated as an ordinal variable reflecting increasing levels of AI deployment maturity. Organizations at more advanced stages of adoption are likely to have developed internal capabilities, infrastructure, and governance mechanisms that lower the cost and risk of further AI integration. Therefore, they are more likely to adopt open-source generative AI tools as they seek flexibility, customization, and cost efficiency.

Perceived Gen AI impacts. To construct a control variable capturing the breadth of perceived organizational benefits from generative AI, we utilized the survey item: “What impacts has generative AI had on your primary use case? (select all that apply).” Respondents could select from the following 10 options: (1) Don’t know or not sure, (2) Too soon to tell, (3) No or negligible impact, (4) Cost reductions, (5) Enhanced innovation, (6) Improved customer satisfaction, (7) Improved decision-making processes, (8) Increased efficiency and productivity, (9) Revenue growth, (10) Other (please specify).

Consistent with our focus on capturing realized organizational value, we excluded items (1) “Don’t know or not sure,” (2) “Too soon to tell,” (3) “No or negligible impact,” and (10) “Other (please specify)” from the analysis. Instead, we focused on the six items (4) “Cost reductions,” (5) “Enhanced innovation,” (6) “Improved customer satisfaction,” (7) “Improved decision-making processes,” (8) “Increased efficiency and productivity,” and (9) “Revenue growth,” which directly reflect substantive impacts of generative AI. Each was treated as a binary indicator (1 if selected, 0 otherwise). We then constructed the variable *Gen AI impacts* as the count of these six impact types selected by a respondent, representing the extent to which generative AI has tangibly benefited the organization across multiple dimensions.

Given that the perceived impacts often correlate with the maturity of AI deployment within an organization, we found this variable to be moderately associated with the organization’s generative AI deployment stage (*Gen AI maturity stage*). To isolate its unique contribution, we applied an orthogonalization procedure between *Perceived Gen AI impacts* and *Gen AI maturity stage*. The residuals from this model, which represent the component of perceived impact breadth uncorrelated with deployment maturity, were used in subsequent models to control for AI value realization independently of stage progression.

Perceived ROI from Gen AI. To account for the influence of perceived return on investment from generative AI adoption, we included a variable labeled Perceived ROI from Gen AI, derived from the item: “How much of your organization’s investment in generative AI has been converted into revenue gain?” Participants responded using a 7-point ordinal scale: (1) “Don’t know or not sure,” (2) “Not applicable: organization has not invested in generative AI,” (3) “Little to no gain,” (4) “Some gain,” (5) “Moderate gain,” (6) “Significant gain,” and (7) “Substantial gain.”

In line with our approach to ensuring interpretability and avoiding noise from ambiguous or inapplicable responses, responses (1) “Don’t know or not sure” is excluded from analysis. The remaining responses (2 through 7) were retained as an ordinal categorical

variable. This variable was included as a control under the theoretical rationale that organizations perceiving higher returns from generative AI are more likely to pursue further cost-saving or innovation-enhancing strategies, such as adopting open-source Gen AI tools. However, because this variable is also empirically correlated with the organization's AI implementation stage (*Gen AI maturity stage*), we applied an orthogonalization procedure to remove any collinearity and isolate its unique variance.

3.2.6. Instrumental Variable

Open GenAI perspectives. To address potential endogeneity concerns arising from the possibility that dependent variable, *Intention to adopt Open GenAI in organization*, and one of the independent variables, *Open-source Orientation* may be jointly influenced by an unobserved factor, we employed an instrumental variable, *Open-source perspectives*. The variable was selected as an instrumental variable because it could capture organizational attitudes toward open-source AI projects, which can be linked to an organization's open-source orientation, but not directly to its intention to adopt Open GenAI project. As Fitzgerald (2006) explains, a favorable perspective toward Open GenAI is not sufficient to drive its implementation [57]. Rather, Open GenAI adoption is facilitated by the maturity of the surrounding ecosystem and infrastructures, which enhance the practical visibility of Open GenAI models. In this context, while a positive view of Open GenAI may be associated with an organization's *Open-source orientation*, it is less likely to serve as the sole determinant of whether the organization intends to adopt Open GenAI. Such conceptual separation supports the validity of our chosen instrument as a plausibly exogenous predictor in addressing endogeneity concerns. Therefore, we considered *Open GenAI perspective* based on the following questions: "To what extent do you agree that open-source AI will. . ." followed by nine potential benefits, including (1) promote equitable distribution of AI benefits across society, (2) accelerate innovation in AI, (3) promote the growth of the AI ecosystem, (4) enhance AI safety and security, (5) enable customization, (6) provide better privacy control, (7) reduce vendor lock-in, (8) promote ethical use of AI, and (9) deliver significant societal and economic benefits. Each item was rated on a five-point Likert scale from "Strongly disagree" to "Strongly agree," with an additional "Don't know or not sure" option that is excluded from analysis. The instrumental variable was constructed by counting the number of benefits for which a respondent selected either "Agree" or "Strongly agree." This variable was used exclusively for endogeneity testing and was not included as a predictor in the main models.

4. Statistical Analysis

To test our hypotheses, we employed ordered logistic regression analysis using the *OrderedModel* class from the *statsmodels* library in Python. This modeling approach was appropriate because the dependent variable is measured on an ordinal scale, reflecting increasing levels of commitment or adoption likelihood. Ordered logistic regression allows us to model the cumulative log-odds of being at or above a particular category of the ordinal outcome, conditional on covariates. The basic form of the ordered logit model can be expressed as:

$$\text{logit}(P(Y \leq j)) = \theta_j - X\beta$$

where Y is the ordinal outcome variable (intention to adopt open-source generative AI), $j \in \{1, 2, \dots, J - 1\}$, denotes thresholds for the J outcome levels, θ_j represents cut-points to be estimated, X is a vector of predictors (e.g., open-source orientation, normative uncertainty, moderating variables), β is the vector of coefficients. This formulation assumes the proportional odds assumption, meaning the effects of covariates are consistent across the cumulative logits.

5. Results

Descriptive statistics for all variables included in the model are summarized in Table 3. This table includes min, max, means, and standard deviations for all key variables. Correlation tables for all variables included in Table 4. Prior to estimating the ordered logistic regression models, we examined the degree of multicollinearity among independent variables using the Variance Inflation Factor (VIF). Across all predictors and interaction terms included in the models, the VIF values ranged from 1.24 to 6.42, which are well below the commonly accepted threshold of 10. These results indicate that multicollinearity is not a significant concern in our model specifications, and the estimated coefficients can be interpreted with confidence. All independent variables used in interaction terms (e.g., normative uncertainty, open-source orientation, moderator variables) were mean-centered prior to estimation to facilitate interpretation and reduce multicollinearity.

Table 3. Descriptive Statistics.

Variable	Min	Max	Mean	SD
DV: Intention to adopt open-source generative AI in organization	1.00	5.00	4.03	0.81
Size	1.00	6.00	3.03	1.86
IT industry	0.00	1.00	0.67	0.47
HQ Location	0.00	1.00	0.47	0.50
Open-source neutral hosting importance	−3.33	3.90	0.10	1.38
Gen AI maturity stage	1.00	5.00	2.93	1.38
Perceived Gen AI impacts	−2.81	4.62	0.22	1.59
Perceived ROI from Gen AI	−2.56	3.25	0.09	1.19
(Hypothesis 1) Open-source Orientation (Mean-centered)	−2.51	1.49	0.55	0.85
Normative Uncertainty (Mean-centered)	−1.10	1.71	−0.04	0.83
(Hypothesis 2) OS Orientation x Normative Uncertainty	−3.45	2.54	−0.08	0.93
Product-oriented application (dummy)	0.00	1.00	0.38	0.49
OS Orientation x Product-oriented	−1.51	1.49	0.23	0.60
Normative Uncertainty x Product-oriented	−1.10	1.71	−0.01	0.52
(Hypothesis 3a) OS Orientation x Normative Uncertainty x Product-oriented	−1.64	2.54	0.01	0.58
Process-oriented application (dummy)	0.00	1.00	0.27	0.44
OS Orientation x Process-oriented	−2.51	1.49	0.16	0.51
Normative Uncertainty x Process-oriented	−1.10	1.71	−0.02	0.45
(Hypothesis 3b) OS Orientation x Normative Uncertainty x Process-oriented	−3.45	2.54	−0.05	0.52

To address potential endogeneity concerns regarding the relationship between dependent variable (*Intention to adopt Open GenAI in organization*) and one of independent variables (*Open-source Orientation*), we conducted a two-stage residual inclusion (2SRI) test using *Open GenAI perspective* as an instrumental variable. The 2SRI test is particularly appropriate for ordered categorical dependent variables as it accommodates the non-linear nature of ordered logistic regression while testing for endogeneity [58]. The first-stage regression demonstrated that the instrumental variable (*Open GenAI perspective*) significantly predicted the endogenous variable *Open-source orientation* (F-statistic = 6.03, $R^2 = 0.194$, $\beta = 0.116$, $p < 0.001$), confirming strong instrument relevance and satisfying the first condition for a valid instrument. In the second stage, we included the first-stage residuals in an ordered logistic regression model predicting *Intention to adopt Open GenAI in organization*. The coefficient of the residuals was not statistically significant ($\beta = 0.687$, $z = 1.231$, $p = 0.218$), indicating that endogeneity is not a concern in our model. The non-significant residuals

term suggests that unobserved factors do not simultaneously influence both *Intention to adopt Open GenAI in organization* and *Open-source Orientation*, supporting the exogeneity assumption of our main predictor. Therefore, we proceeded with standard ordered logistic regression analysis without instrumental variable correction.

Table 4. Correlation Table.

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)
(1) Intention to adopt open-source generative AI in organization	1.00	0.02	0.11	−0.16	0.06	0.10	0.14	0.10	0.25	0.10	−0.01	0.06	0.23	0.07	0.11	0.01	0.13	0.14	−0.06
(2) Size	0.02	1.00	−0.17	0.12	0.01	0.09	0.10	0.13	−0.19	0.28	0.12	−0.03	−0.12	0.14	0.03	0.00	−0.12	0.14	0.13
(3) IT Industry	0.11	−0.17	1.00	−0.04	0.01	0.14	0.18	0.16	0.18	−0.09	−0.01	0.07	0.05	−0.22	−0.14	0.04	0.10	0.04	0.04
(4) HQ Location	−0.16	0.12	−0.04	1.00	0.02	0.06	−0.11	−0.12	−0.11	0.07	0.04	−0.09	−0.11	0.04	0.05	0.12	−0.03	−0.03	−0.02
(5) Open-source Neutral Hosting Importance	0.06	0.01	0.01	0.02	1.00	−0.08	0.20	0.15	−0.08	0.13	0.01	0.00	0.06	0.06	0.03	−0.06	−0.30	0.18	−0.03
(6) Gen AI maturity stage	0.10	0.09	0.14	0.06	−0.08	1.00	0.05	0.04	0.14	0.08	0.08	0.02	0.13	0.03	0.12	0.24	0.08	0.06	−0.05
(7) Perceived Gen AI Impacts	0.14	0.10	0.18	−0.11	0.20	0.05	1.00	0.40	0.13	0.13	0.06	0.07	0.11	0.02	−0.03	0.03	0.05	0.15	0.18
(8) ROI from Gen AI	0.10	0.13	0.16	−0.12	0.15	0.04	0.40	1.00	0.07	−0.07	−0.03	−0.03	−0.01	−0.02	−0.02	0.03	0.08	−0.04	0.03
(9) Open-source (OS) Orientation	0.25	−0.19	0.18	−0.11	−0.08	0.14	0.13	0.07	1.00	−0.08	0.11	0.06	0.57	0.03	0.05	0.05	0.45	−0.10	0.11
(10) Normative Uncertainty	0.10	0.28	−0.09	0.07	0.13	0.08	0.13	−0.07	−0.08	1.00	0.55	0.02	0.03	0.62	0.41	−0.03	−0.10	0.53	0.22
(11) OS Orientation × Normative Uncertainty	−0.01	0.12	−0.01	0.04	0.01	0.08	0.06	−0.03	0.11	0.55	1.00	0.08	0.08	0.41	0.63	−0.07	0.07	0.23	0.56
(12) Process-oriented Application of Gen AI	0.06	−0.03	0.07	−0.09	0.00	0.02	0.07	−0.03	0.06	0.02	0.08	1.00	0.50	−0.02	0.02	−0.47	−0.25	0.04	0.08
(13) OS Orientation × Process-oriented Application	0.23	−0.12	0.05	−0.11	0.06	0.13	0.11	−0.01	0.57	0.03	0.08	0.50	1.00	0.03	0.07	−0.23	−0.13	0.02	0.04
(14) Normative Uncertainty × Process-oriented Application	0.07	0.14	−0.22	0.04	0.06	0.03	0.02	−0.02	0.03	0.62	0.41	−0.02	0.03	1.00	0.66	0.01	0.01	0.00	0.00
(15) OS Orientation × Normative Uncertainty × Process-oriented Application	0.11	0.03	−0.14	0.05	0.03	0.12	−0.03	−0.02	0.05	0.41	0.63	0.02	0.07	0.66	1.00	−0.01	0.00	0.00	0.00
(16) Product-oriented Application of Gen AI	0.01	0.00	0.04	0.12	−0.06	0.24	0.03	0.03	0.05	−0.03	−0.07	−0.47	−0.23	0.01	−0.01	1.00	0.54	−0.09	−0.16
(17) OS Orientation × Product-oriented Application	0.13	−0.12	0.10	−0.03	−0.30	0.08	0.05	0.08	0.45	−0.10	0.07	−0.25	−0.13	0.01	0.00	0.54	1.00	−0.21	0.10
(18) Normative Uncertainty × Product-oriented Application	0.14	0.14	0.04	−0.03	0.18	0.06	0.15	−0.04	−0.10	0.53	0.23	0.04	0.02	0.00	0.00	−0.09	−0.21	1.00	0.41
(19) OS Orientation × Normative Uncertainty × Product-oriented Application	−0.06	0.13	0.04	−0.02	−0.03	−0.05	0.18	0.03	0.11	0.22	0.56	0.08	0.04	0.00	0.00	−0.16	0.10	0.41	1.00

Table 5 reports the results of the ordered logistic regression models, which examined organizational intention to adopt open-source generative AI (Gen AI). Model 1 includes only control variables, such as firm size, industry sector, HQ location, and Gen AI maturity stage metrics. While not all are statistically significant, larger firms (size) and those in the IT industry (tend to show higher intentions to adopt Open GenAI. Similarly, organizations with greater ROI from Gen AI, more advanced Gen AI maturity in organizations, and stronger perceived Gen AI impacts on business also demonstrate increased adoption intentions. These results suggest that organizations with greater familiarity and maturity in Gen AI implementation tend to be more willing and ready to embrace Open GenAI solutions. Model 2 introduces the main effect of Open-source Orientation. As hypothesized (H1), a positive and statistically significant association is observed ($\beta = 0.57$, $p < 0.01$), indicating that organizations with stronger open-source values are more likely to adopt open-source Gen AI. This finding aligns with our theoretical expectation that organizations possessing deep knowledge and well-established routines around open-source demonstrate greater absorptive and desorptive capacities, which in turn enhance their intention to adopt Open GenAI technologies. Model 3 introduces the interaction between open-source

orientation and normative uncertainty to test Hypothesis 2. While open-source orientation shows a consistently positive and significant effect on Open GenAI adoption ($\beta = 0.61$, $p < 0.001$ in Model 3), the interaction term with normative uncertainty is significantly negative ($\beta = -0.36$, $p < 0.05$). This suggests that under conditions of elevated ethical or regulatory concerns, even organizations with strong open-source values may exhibit lower adoption intentions, which supports H2. To examine the moderating role of application context, Model 4 introduces Process-oriented Applications and their higher-order interaction with open-source orientation and normative uncertainty. The three-way interaction term (open-source orientation $\times$ normative uncertainty $\times$ process-oriented Application) is positive and statistically significant ($\beta = 1.26$, $p < 0.01$), providing support for Hypothesis 3a. This indicates that when open-source generative AI is applied to internal processes, such as workflow automation or internal analytics, the weakening effect of normative uncertainty is mitigated.

Table 5. Ordered Logit Analysis.

Variable	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
Size	0.03 (0.08)	0.09 (0.08)	0.06 (0.08)	0.10 (0.08)	0.11 (0.08)	0.14 (0.08)
IT Industry	0.27 (0.30)	0.22 (0.30)	0.25 (0.30)	0.37 (0.31)	0.23 (0.30)	0.42 (0.32)
HQ Location	−0.50 (0.28)	−0.50 (0.28)	−0.49 (0.28)	−0.51 (0.28)	−0.47 (0.28)	−0.45 (0.29)
Open-source Neutral Hosting Importance	0.12 (0.10)	0.14 (0.10)	0.12 (0.10)	0.09 (0.10)	0.13 (0.11)	0.15 (0.11)
Gen AI maturity stage	0.14 (0.09)	0.11 (0.10)	0.08 (0.10)	0.11 (0.10)	0.09 (0.10)	0.09 (0.10)
Perceived Gen AI Impacts	0.09 (0.11)	0.06 (0.11)	0.07 (0.11)	0.00 (0.11)	0.05 (0.12)	0.02 (0.12)
ROI from Gen AI	0.08 (0.13)	0.07 (0.13)	0.11 (0.13)	0.13 (0.13)	0.12 (0.13)	0.13 (0.13)
H1) Open-source (OS) Orientation		0.57 ** (0.18)	0.61 ** (0.18)	0.64 *** (0.18)	0.73 *** (0.18)	0.67 *** (0.18)
Normative Uncertainty			0.43 (0.21)	0.38 (0.21)	0.31 (0.22)	0.33 (0.22)
H2) OS Orientation $\times$ Normative Uncertainty			−0.36 * (0.18)	−0.39 * (0.18)	−0.27 (0.19)	−0.42 * (0.19)
Process-oriented Application of Gen AI				0.04 (0.34)		−0.10 (0.38)
OS Orientation $\times$ Process-oriented Application				0.37 (0.34)		0.80 * (0.38)
Normative Uncertainty $\times$ Process-oriented Application				−0.58 (0.43)		−0.24 (0.49)
H3a) OS Orientation $\times$ Normative Uncertainty $\times$ Process-oriented Application				1.26 ** (0.39)		1.19 ** (0.46)
Product-oriented Application of Gen AI					−0.53 (0.41)	−0.57 (0.46)
OS Orientation $\times$ Product-oriented Application					0.62 (0.42)	1.14 * (0.47)
Normative Uncertainty $\times$ Product-oriented Application					0.92 * (0.44)	0.85 (0.50)
H3b) OS Orientation $\times$ Normative Uncertainty $\times$ Product-oriented Application					−0.94 * (0.39)	−0.30 (0.46)
Sample Size	209	209	209	209	209	209
Log-Likelihood	−227.90	−222.64	−220.11	−213.43	−216.28	−209.55
AIC	477.80	469.30	468.20	462.90	468.60	463.10
BIC	514.60	509.40	515.00	523.00	528.70	536.60

Notes: $p < 0.1$; * $p < 0.05$; ** $p < 0.01$; *** $p < 0.001$.

In other words, open-source oriented organizations are more resilient to uncertainty when open-source generative AI deployment is applied within organizational boundaries. Model 5 shifts the focus to product-oriented applications to test Hypothesis 3b. Here, the three-way interaction (open-source orientation $\times$ normative uncertainty $\times$ product-oriented Application) is negative and marginally significant in Model 5 ($\beta = -0.94, p < 0.05$) but loses significance in the full model (Model 6: $\beta = -0.30$, not supported). This result offers partial support for H3b, suggesting that when open-source generative AI is embedded in externally facing products or services, normative concerns may amplify and suppress the influence of open-source orientation on adoption intentions. Thus, the visibility and accountability associated with external deployments appear to intensify the dampening effect of normative risk. Finally, Model 6 incorporates all interaction terms, serving as the full model. The interaction between open-source orientation and normative uncertainty remains significant ($\beta = -0.73, p < 0.05$), reaffirming H2. The three-way interaction for process-oriented use cases (H3a) continues to hold ($\beta = 1.19, p < 0.01$), while the product-oriented three-way term (H3b) is not statistically significant. This reduction in significance for H3b is likely due to several interconnected factors. First, the process- and product-oriented application variables originate from a single-choice survey question, which inherently produces moderate negative correlation between these variables ($r = -0.47$). When both are included simultaneously along with their interaction terms, especially higher-order interactions, this correlation may inflate standard errors and reduce the power to detect significant effects. Second, the H3a and H3b terms explain overlapping variance in the intention to adopt Open GenAI, meaning they partially compete in accounting for the same variance. This shared explanatory power diminishes the unique contribution of the H3b term when both are modeled together. Despite this, Model 6 remains valuable as it captures the joint moderating roles of different application contexts and provides a nuanced understanding of how normative uncertainty influences the relationship between open-source orientation and adoption intentions.

6. Discussion

This study examined how open-source oriented organizations adopt open-source generative AI under conditions of normative uncertainty and how the application contexts of internal processes or of external products and services moderate this relationship. Drawing on organizational learning theory and open innovation literature, we theorized that organizations with strong open-source orientation would generally be more favorable to open-source generative AI adoption. However, we also found that this relationship is weakened by the degree of normative uncertainty perceived. Furthermore, we propose that the context in which open-source generative AI is applied, specifically, whether it is integrated into internal organizational processes or external-facing products, shapes the extent to which normative uncertainty influences adoption decisions. Consistent with our expectations, the results indicate that open-source oriented organizations show greater adoption intentions for open-source generative AI, but when normative uncertainty is high, the positive impact of such orientation diminishes. Moreover, our results confirm that the moderating effect of normative uncertainty differs by open-source generative AI application context. The impact of normative uncertainty is weaker when open-source generative AI is applied to internal processes but becomes stronger and more deterrent when used in external-facing products. This underscores the importance of considering not only organizational capabilities but also the visibility of deployment and associated reputational risks in navigating normative uncertainty.

6.1. Theoretical Implications

This study contributes to organizational research by integrating insights from organizational learnings, open innovation, and normative uncertainty. First, we extend the absorptive and desorptive capacity frameworks by highlighting the conditional nature of open innovative knowledge absorption under normatively uncertain environments. While open-source oriented organizations tend to exhibit stronger capabilities for adopting community-driven technologies like open-source generative AI, our findings reveal that normative uncertainty, such as concerns over privacy, bias, or misuse, can disrupt the learning routines associated with these capacities. This underscores normative risk as a critical but underexplored boundary condition in the effective functioning of absorptive capacity [4,5]. Second, by distinguishing between exploitative (internal processes) and exploratory (external products or services) applications of open-source generative AI, the study contributes to the literature on organizational learning modes [4]. We show that normative uncertainty exerts a differentiated effect depending on application context. While internal deployments allow organizations to manage risks through bounded control, the integration of open-source generative AI into products or services exposes organizations to reputational and regulatory pressures, thereby amplifying the salience of normative concerns. This context-dependence extends dual-process innovation theories [14] and suggests that the impact of normative uncertainty is contingent not only on organizational capabilities but also on where and how the technology is applied. Third, the study deepens our understanding of open innovation under conditions of normative uncertainty. Although prior research has celebrated the democratizing and collaborative advantages of openness [7,9], less attention has been paid to how legitimacy pressures may reduce the adoption of open innovation, even within open-source oriented organizations. By showing that normative uncertainty can override internal readiness, our findings call for an expansion of open innovation frameworks to explicitly incorporate ethical and political dimensions. This finding contributes to the Technology Acceptance Model (TAM), which traditionally focuses on Perceived Usefulness and Perceived Ease of Use [59]. Extending TAM beyond social influence [60], our study suggests that an organization's perceived legitimacy in adopting Open GenAI, which is derived from collective ethical concerns, significantly shapes the decision-making process. Furthermore, expanding our discussion beyond the organizational level, it would also be fruitful to explore how normative uncertainty surrounding a disruptive technology can influence an entire industry's adoption trajectory. A recent study by Ambrozio, Lindeque, and Peter (2025) on cloud computing, for instance, found that as uncertainty decreases, firms shift their adoption drivers from institutional isomorphism pressure to competitive isomorphism pressure [61,62]. However, our research on GenAI adoption shows a different context. With emerging regulations and ongoing normative uncertainty around ethical use [48], organizations might instead be primarily influenced by mimetic isomorphism pressure. This suggests that in the early stages of a technology's life cycle, when norms are still being established, organizations may adopt GenAI not for competitive advantage or institutional mandates, but simply by imitating other organizations to reduce risk and to acquire legitimacy.

6.2. Practical Implications

From a practical standpoint, this study suggests that for risk-contained experimentation and learning, organizations should adopt a staged approach to open-source generative AI deployment by starting with internal, process-oriented applications where ethical scrutiny is limited. As regulatory frameworks like the EU AI Act evolve, legal and compliance teams must stay engaged with open-source licensing, data provenance, and model auditability. Forward-looking organizations may proactively contribute to shaping these

standards through participation in multi-stakeholder forums. Second, our findings emphasize the importance of cultivating normative readiness in parallel, such as building internal ethics review boards, transparent documentation practices, and active communication between stakeholders to manage normative ambiguity. For instance, recently many organizations have considered implementing an Open-Source Program Office (OSPO), which is a designated team or individual within an organization responsible for defining, coordinating, and managing the organization's open-source strategy to provide guidance on legal compliance and best practices. From a normative uncertainty perspective on Open GenAI adoption, this function offers significant practical benefits. Given the ethical ambiguities on Open GenAI from data privacy to intellectual property, the implementation of an OSPO can reduce normative uncertainty by establishing clear internal guidelines and a code of conduct. Lastly, these findings are also important for regulators and policymakers. To support responsible innovation, clearer and more consistent rules, especially around licensing, data sources, and how AI models are checked, are needed to reduce any confusion on normative uncertainty in Open GenAI usage. Providing such guidelines could clarify expectations and acceptable practices, thereby mitigating the effects of normative uncertainty on Open GenAI adoption. For instance, according to the U.S. AI Action Plan (2025) [63], open-source and open-weight AI models enable startups to innovate without dependence on closed providers, allow governments and businesses to work with sensitive data securely, and are essentials for advancing academic research. Embedding these principles into policy guidance can implement Open GenAI solutions responsibly.

6.3. Limitations and Future Research

Despite its contributions, this study has several limitations that open avenues for future research. First, our cross-sectional design limits causal inference. Collected from a Linux Foundation survey, our dataset likely overrepresents organizations that are already supportive of open-source technologies. Since the survey only categorized respondents as IT or non-IT without capturing detailed industry differences, our findings may not fully reflect the challenges and opportunities that diverse sectors, such as finance or healthcare, face when adopting open-source GenAI. In addition, key organizational characteristics such as age were not asked, and size was captured only via Likert scales, which limits the granularity of demographic insights. Like all survey-based research, our study only captures organizations' intentions to adopt, which may not always translate into actual implementation. Considering these empirical limitations, future research would benefit from more diverse industry samples and longitudinal designs to track the progression from intended adoption to actual use of Open GenAI. Moreover, incorporating organizational demographic factors from an ecological perspective [64] could provide richer understanding of how structural characteristics influence adoption patterns. Such longitudinal studies could also provide deeper insights into how normative uncertainty changes throughout Open GenAI adoption stages from initial experimentation to full deployment, and how it interacts with organizational learning processes over time [5]. Second, although we differentiate between internal and external application contexts, we do not examine how the scope of open-source integration within organizations, such as depth across functions or roles, might shape perceptions of normative uncertainty. Future work could explore whether wider internal exposure to open-source tools moderates ethical concerns or builds normative resilience. For instance, building on TAM's emphasis on perceived usefulness and ease of use, future research could investigate how varying levels of open-source integration impacts employees' attitudes and behavioral intentions towards Open GenAI adoption or whether wider internal exposure to open-source tools moderates ethical concerns or builds normative resilience within an organization [60]. Moreover, examining

how organizational learning and socialization processes mediate the relationship between normative uncertainty and technology acceptance would provide deeper insights into how responsible GenAI to be used [65]. Third, although this study focuses on organizations adopting open-source generative AI, the future study can explore the different scenarios that organizations replace proprietary AI systems with open-source generative AI due to cost or innovation needs, the domain of which may be underexplored. These transitions may entail environmental shocks that trigger legitimacy reassessment, governance restructuring, or reputational repositioning [66,67]. Due to limitations in the current data, our study only focuses exclusively on the adoption of Open GenAI. However, future research should examine whether normative uncertainty affects the adoption of generative AI overall or primarily encourages organizations to pursue proprietary AI alternatives instead. This distinction is important to better understand how normative uncertainty shapes AI adoption strategies more broadly. Fourth, we did not examine the impact of network positions within open-source communities in the context of open innovation, where organizations' positions, such as centrality, structural holes [68], or status position in the communities [69], may influence both the quality of information exchanged and the framing of normative concerns. For instance, high-status contributors in open innovation communities may perceive lower normative risk when they adopt open-source generative AI, as their established legitimacy and centrality afford them interpretive authority and reduce the likelihood of external contestation. As Kim [70] shows, status buffers organizations from reputational penalties in uncertain institutional environments when they adopt normatively uncertain innovation. Finally, future research could refine the operationalization of normative uncertainty beyond ethical and/or regulatory perception items. Multidimensional scales that distinguish fairness, explainability, privacy, or accountability could yield greater analytical precision [71,72]. As generative AI matures and open models proliferate, future inquiries should continue to examine how organizations interpret and respond to evolving norms, especially in open, decentralized innovation environments.

Author Contributions: Conceptualization, S.H. and D.Y.; methodology, S.H., H.R., and X.J.; software, S.H., H.R., and; validation, S.H., H.R. and X.J.; formal analysis, S.H., H.R., X.J., and D.Y.; investigation, S.H.; resources, S.H.; data curation, S.H., H.R., and X.J.; writing—original draft preparation, S.H., H.R., and X.J.; writing—review and editing, S.H. and D.Y.; supervision, S.H. and D.Y.; project administration, D.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.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data can be obtained on request at The Linux Foundation (<https://www.linuxfoundation.org/research> (accessed on 1 May 2025)).

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

References

1. Holmström, J.; Carroll, N. How organizations can innovate with generative AI. *Bus. Horiz.* **2025**, *68*, 559–573. [\[CrossRef\]](#)
2. Mariani, M.; Dwivedi, Y.K. Generative artificial intelligence in innovation management: A preview of future research developments. *J. Bus. Res.* **2024**, *175*, 114542. [\[CrossRef\]](#)
3. López-Solís, O.; Luzuriaga-Jaramillo, A.; Bedoya-Jara, M.; Naranjo-Santamaría, J.; Bonilla-Jurado, D.; Acosta-Vargas, P. Effect of generative artificial intelligence on strategic decision-making in entrepreneurial business initiatives: A systematic literature review. *Adm. Sci.* **2025**, *15*, 66. [\[CrossRef\]](#)
4. Cohen, W.M.; Levinthal, D.A. Innovation and Learning: The Two Faces of R & D. *Econ. J.* **1989**, *99*, 569–596. [\[CrossRef\]](#)
5. Zahra, S.A.; George, G. Absorptive capacity: A review, reconceptualization, and extension. *Acad. Manag. Rev.* **2002**, *27*, 185–203. [\[CrossRef\]](#)

6. Chesbrough, H.W. *Open Innovation: The New Imperative for Creating and Profiting from Technology*; Harvard Business Press: Brighton, MA, USA, 2003.
7. West, J.; Gallagher, S. Challenges of open innovation: The paradox of firm investment in open-source software. *Rd Manag.* **2006**, *36*, 319–331. [[CrossRef](#)]
8. Baldwin, C.; Von Hippel, E. Modeling a paradigm shift: From producer innovation to user and open collaborative innovation. *Organ. Sci.* **2011**, *22*, 1399–1417. [[CrossRef](#)]
9. Von Krogh, G.; Von Hippel, E. The promise of research on open source software. *Manag. Sci.* **2006**, *52*, 975–983. [[CrossRef](#)]
10. Lakhani, K.; Panetta, J. The principles of distributed innovation. In *Successful OSS Project Design and Implementation*; Routledge: London, UK, 2016; pp. 7–26.
11. March, J.G. Exploration and exploitation in organizational learning. *Organ. Sci.* **1991**, *2*, 71–87. [[CrossRef](#)]
12. Galleli, B.; Amaral, L. Bridging Institutional Theory and Social and Environmental Efforts in Management: A Review and Research Agenda. *J. Manag.* **2025**. [[CrossRef](#)]
13. Rhee, M.; Valdez, M.E. Contextual factors surrounding reputation damage with potential implications for reputation repair. *Acad. Manag. Rev.* **2009**, *34*, 146–168. [[CrossRef](#)]
14. Tushman, M.L.; O'Reilly, C.A., III. Ambidextrous organizations: Managing evolutionary and revolutionary change. *Calif. Manag. Rev.* **1996**, *38*, 8–29. [[CrossRef](#)]
15. Benders, J.; Van Veen, K. What's in a fashion? Interpretative viability and management fashions. *Organization* **2001**, *8*, 33–53. [[CrossRef](#)]
16. Fabrizio, K.R. Absorptive capacity and the search for innovation. *Res. Policy* **2009**, *38*, 255–267. [[CrossRef](#)]
17. Lane, P.J.; Koka, B.R.; Pathak, S. The reification of absorptive capacity: A critical review and rejuvenation of the construct. *Acad. Manag. Rev.* **2006**, *31*, 833–863. [[CrossRef](#)]
18. Lichtenthaler, U.; Lichtenthaler, E. A capability-based framework for open innovation: Complementing absorptive capacity. *J. Manag. Stud.* **2009**, *46*, 1315–1338. [[CrossRef](#)]
19. Aliasghar, O.; Haar, J. Open innovation: Are absorptive and desorptive capabilities complementary? *Int. Bus. Rev.* **2023**, *32*, 101865. [[CrossRef](#)]
20. Chesbrough, H.; Lettl, C.; Ritter, T. Value creation and value capture in open innovation. *J. Prod. Innov. Manag.* **2018**, *35*, 930–938. [[CrossRef](#)]
21. Kim, E.; Lee, I.; Kim, H.; Shin, K. Factors affecting outbound open innovation performance in bio-pharmaceutical industry-focus on out-licensing deals. *Sustainability* **2021**, *13*, 4122. [[CrossRef](#)]
22. Hu, Y.; McNamara, P.; McLoughlin, D. Outbound open innovation in bio-pharmaceutical out-licensing. *Technovation* **2015**, *35*, 46–58. [[CrossRef](#)]
23. Chesbrough, H. Explicating open innovation: Clarifying an emerging paradigm for understanding innovation. In *New Frontiers in Open Innovation*; Chesbrough, H., Vanhaverbeke, W., West, J., Eds.; Oxford University Press: Oxford, UK, 2014; pp. 3–28.
24. Lane, P.J.; Lubatkin, M. Relative absorptive capacity and interorganizational learning. *Strateg. Manag. J.* **1998**, *19*, 461–477. [[CrossRef](#)]
25. Chiaroni, D.; Chiesa, V.; Frattini, F. Unravelling the process from closed to open innovation: Evidence from mature, asset-intensive industries. *Rd Manag.* **2010**, *40*, 222–245. [[CrossRef](#)]
26. Scott, W.R. *Institutions and Organizations*, 2nd ed.; Sage Publications: Thousand Oaks, CA, USA, 2001.
27. Suchman, M.C. Managing legitimacy: Strategic and institutional approaches. *Acad. Manag. Rev.* **1995**, *20*, 571–610. [[CrossRef](#)]
28. Ansari, S.M.; Fiss, P.C.; Zajac, E.J. Made to fit: How practices vary as they diffuse. *Acad. Manag. Rev.* **2010**, *35*, 67–92.
29. Kim, T.; Yang, D. Multiple Goals, Attention Allocation, and the Intention-Achievement Gap in Energy Efficiency Innovation. *Sustainability* **2020**, *12*, 7102. [[CrossRef](#)]
30. Zietsma, C.; Lawrence, T.B. Institutional work in the transformation of an organizational field: The interplay of boundary work and practice work. *Adm. Sci. Q.* **2010**, *55*, 189–221. [[CrossRef](#)]
31. March, J.G.; Olsen, J.P. The uncertainty of the past: Organizational learning under ambiguity. *Eur. J. Political Res.* **1975**, *3*, 147–171. [[CrossRef](#)]
32. Jalonon, H. The uncertainty of innovation: A systematic review of the literature. *J. Manag. Res.* **2012**, *4*, 1–47. [[CrossRef](#)]
33. Rivette, K.G.; Kline, D. *Rembrandts in the Attic: Unlocking the Hidden Value of Patents*; Harvard Business School Press: Boston, MA, USA, 2000.
34. Fosfuri, A. The licensing dilemma: Understanding the determinants of the rate of technology licensing. *Strateg. Manag. J.* **2006**, *27*, 1141–1158. [[CrossRef](#)]
35. Lichtenthaler, U. The drivers of technology licensing: An industry comparison. *Calif. Manag. Rev.* **2007**, *49*, 67–89. [[CrossRef](#)]
36. Deephouse, D.L.; Carter, S.M. An examination of differences between organizational legitimacy and organizational reputation. *J. Manag. Stud.* **2005**, *42*, 329–360. [[CrossRef](#)]

37. Mishina, Y.; Block, E.S.; Mannor, M.J. The path dependence of organizational reputation: How social judgment influences assessments of capability and character. *Acad. Manag. J.* **2012**, *55*, 971–995. [\[CrossRef\]](#)
38. Pfarrer, M.D.; Decelles, K.A.; Smith, K.G.; Taylor, M.S. After the fall: Reintegrating the corrupt organization. *Acad. Manag. Rev.* **2008**, *33*, 730–749. [\[CrossRef\]](#)
39. Sims, R.R. Toward a better understanding of organizational efforts to rebuild reputation following an ethical scandal. *J. Bus. Ethics* **2009**, *90*, 453–472. [\[CrossRef\]](#)
40. Gupta, A.K.; Smith, K.G.; Shalley, C.E. The interplay between exploration and exploitation. *Acad. Manag. J.* **2006**, *49*, 693–706. [\[CrossRef\]](#)
41. Lavie, D.; Stettner, U.; Tushman, M.L. Exploration and exploitation within and across organizations. *Acad. Manag. Ann.* **2010**, *4*, 109–155. [\[CrossRef\]](#)
42. March, J.G. Rationality, foolishness, and adaptive intelligence. *Strateg. Manag. J.* **2006**, *27*, 201–214. [\[CrossRef\]](#)
43. Bommasani, R.; Hudson, D.A.; Adeli, E.; Altman, R.; Arora, S.; von Arx, S.; Bernstein, M.S.; Bohg, J.; Bosselut, A.; Brunskill, E. On the opportunities and risks of foundation models. *arXiv* **2021**, arXiv:2108.07258. [\[CrossRef\]](#)
44. Microsoft. What Is Responsible AI? Available online: <https://learn.microsoft.com/en-us/azure/machine-learning/concept-responsible-ai?view=azureml-api-2> (accessed on 29 June 2025).
45. Fjeld, J.; Achten, N.; Hilligoss, H.; Nagy, A.; Srikumar, M. *Principled Artificial Intelligence*; Berkman Klein Center: Cambridge, MA, USA, 2020; Available online: <https://cyber.harvard.edu/publication/2020/principled-ai> (accessed on 29 June 2025).
46. Dahlander, L.; Gann, D.M. How open is innovation? *Res. Policy* **2010**, *39*, 699–709. [\[CrossRef\]](#)
47. IBM. *Securing Generative AI: Risk Taxonomy and Mitigation Playbook*; IBM Security Report: Armonk, NY, USA, 2024.
48. Hacker, P. Sustainable AI regulation. *Common Mark. Law Rev.* **2024**, *61*, 345–386. [\[CrossRef\]](#)
49. Tschang, F.T.; Almirall, E. Artificial intelligence as augmenting automation: Implications for employment. *Acad. Manag. Perspect.* **2021**, *35*, 642–659. [\[CrossRef\]](#)
50. Raisch, S.; Fomina, K. Combining human and artificial intelligence: Hybrid problem-solving in organizations. *Acad. Manag. Rev.* **2025**, *50*, 441–464. [\[CrossRef\]](#)
51. Ramanujam, R.; Goodman, P.S. Latent errors and adverse organizational consequences: A conceptualization. *J. Organ. Behav. Int. J. Ind. Occup. Organ. Psychol. Behav.* **2003**, *24*, 815–836. [\[CrossRef\]](#)
52. Dobbe, R.; Gilbert, T.K.; Mintz, Y. Hard choices in artificial intelligence: Addressing normative uncertainty through sociotechnical commitments. *arXiv* **2019**, arXiv:1911.09005. [\[CrossRef\]](#)
53. Schlaile, M.P.; Urmetzer, S.; Blok, V.; Andersen, A.D.; Timmermans, J.; Mueller, M.; Fagerberg, J.; Pyka, A. Innovation systems for transformations towards sustainability? Taking the normative dimension seriously. *Sustainability* **2017**, *9*, 2253. [\[CrossRef\]](#)
54. Bartholomew, D.J.; Steele, F.; Moustaki, I. *Analysis of Multivariate Social Science Data*; CRC Press: Boca Raton, FL, USA, 2008.
55. Olsson, U. Maximum likelihood estimation of the polychoric correlation coefficient. *Psychometrika* **1979**, *44*, 443–460. [\[CrossRef\]](#)
56. West, J.; O'mahony, S. The role of participation architecture in growing sponsored open source communities. *Ind. Innov.* **2008**, *15*, 145–168. [\[CrossRef\]](#)
57. Fitzgerald, B. The transformation of open source software. *MIS Q.* **2006**, *30*, 587–598. [\[CrossRef\]](#)
58. Terza, J.V.; Basu, A.; Rathouz, P.J. Two-stage residual inclusion estimation: Addressing endogeneity in health econometric modeling. *J. Health Econ.* **2008**, *27*, 531–543. [\[CrossRef\]](#) [\[PubMed\]](#)
59. Davis, F.D. Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Q.* **1989**, *13*, 319–340. [\[CrossRef\]](#)
60. Venkatesh, V.; Davis, F.D. A theoretical extension of the technology acceptance model: Four longitudinal field studies. *Manag. Sci.* **2000**, *46*, 186–204. [\[CrossRef\]](#)
61. Ambrozio, S.; Lindeque, J.P.; Peter, M.K. Navigating uncertainty: Isomorphic pressures in cloud computing adoption. In *The Palgrave Handbook of Breakthrough Technologies in Contemporary Organisations*; Moussa, M., McMurray, A., Eds.; Springer Nature: Singapore, 2025; Chapter 18. [\[CrossRef\]](#)
62. DiMaggio, P.J.; Powell, W.W. The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *Am. Sociol. Rev.* **1983**, *48*, 147–160. [\[CrossRef\]](#)
63. United States White House. Winning the Race: America's AI Action Plan. 2025. Available online: <https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf> (accessed on 14 August 2025).
64. Hannan, M.T.; Freeman, J. The population ecology of organizations. *Am. J. Sociol.* **1977**, *82*, 929–964. [\[CrossRef\]](#)
65. Weick, K.E. *Sensemaking in Organizations*; Sage Publications: Thousand Oaks, CA, USA, 1995.
66. Tushman, M.L.; Anderson, P. Technological discontinuities and organizational environments. In *Organizational Innovation*; Routledge: Abingdon, UK, 2018; pp. 345–372.
67. Ansell, C.; Boin, A. Taming deep uncertainty: The potential of pragmatist principles for understanding and improving strategic crisis management. *Adm. Soc.* **2019**, *51*, 1079–1112. [\[CrossRef\]](#)
68. Burt, R.S. *Structural Holes: The Social Structure of Competition*; Harvard University Press: Cambridge, MA, USA, 1992.

69. Podolny, J.M. *Status Signals: A Sociological Study of Market Competition*; Princeton University Press: Princeton, NJ, USA, 2010.
70. Kim, B.K. Normative uncertainty and middlestatus innovation in the US daily newspaper industry. *Strateg. Organ.* **2019**, *18*, 377–406. [[CrossRef](#)]
71. Mittelstadt, B.D.; Allo, P.; Taddeo, M.; Wachter, S.; Floridi, L. The ethics of algorithms: Mapping the debate. *Big Data Soc.* **2016**, *3*, 2053951716679679. [[CrossRef](#)]
72. Martin, K. Ethical implications and accountability of algorithms. *J. Bus. Ethics* **2019**, *160*, 835–850. [[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.