

A digital design framework for algorithmic mediation with brands

Bonds in the Age of Intelligent Products





Hola, Hello, こんにちは, Hallo, 你好呀, Olá, Bonjour, Ola!

making tomorrow tangible

- 00 Context
- 01 Hypothesis
- 02 Market Trends
- 03 Sector Benchmark
- 04 Possible Scenarios
- 05 Design Principles
- 06 Takeaways
- 07 Sources

Context

A generation of bilateral relationship between brand and user

For a generation, the main channel through which a brand relates to its customers has been an interface designed, built, and controlled by the brand itself: the website, the app, the product page. The user searched, compared across tabs, and decided. **The work of product design consisted of making that space as clear, persuasive, and fluid as possible for a person who reads, navigates, and clicks.**

That model had already been evolving. Search engines began summarizing instead of just linking, recommendation engines personalized results, comparison sites aggregated third-party offers. But at every one of those steps, **it was still a person who navigated and decided.** The relationship, though mediated by technology, kept its bilateral logic: **brand and user, connected through an interface designed for humans.**

What changes now is not one more channel. It's who holds the interface. A growing share of the work of discovering, comparing, and deciding is being delegated to a personal AI agent that acts on the user's behalf, across multiple brands, without the user needing (or wanting) to see each site or app separately. **Today only 19% of consumers use AI agents to interact with brands, but that figure is expected to reach 46% by the end of 2026.** It's a decade's worth of change compressed into one year.

Two movements converge in the same year

Two movements are happening at the same time, and it's that simultaneity that turns 2026 into an inflection point rather than a future prediction.

The first is behavioral: 45% of online shoppers in the US already use AI tools as part of their purchasing process, and of those, 39% try brands they wouldn't have considered without that assistance. The user no longer reaches the brand the way they did two years ago.

The second is infrastructural: the technical layer behind this shift is being built in parallel by every major AI provider, and a new software category has already emerged just to measure and manage a brand's visibility inside these systems, in a context where Google AI Mode and Perplexity cite sources very differently than ChatGPT does.

Behavior and infrastructure rarely move at the same speed. When they do, the window for a discipline, in this case product design, to define how it responds to that change is short. That is the ground this paper wants to occupy.

Hypothesis

We are transitioning from a direct relationship between brands and users to one mediated by an AI agent. This third actor **curates the options from which the user makes the final choice**

We analyze this new relationship model through six different lenses:

01

Power dynamics

Control shifts from the brand's ecosystem toward a small number of AI infrastructure providers. **Google AI Mode and Perplexity draw around 90% of their citations from the top 10 search results. ChatGPT, by contrast, only 30%.** That divergence already sustains an entire software category, GEO/AEO, dedicated to managing a brand's visibility to these systems. Microsoft, IBM, Nvidia, and Anthropic lead that infrastructure market today.

02

Value exchange

For the user, that shift already pays off, as Context showed. For the brand, the value is harder to win and more fragile than expected. **A study that causally evaluated how agents from different providers shop found broad, model-specific position effects, and found that minor edits to a product listing can substantially shift its selection probability.** Algorithmic preference is the new objective to win, and also a manipulable one.

03

User agency

In the final decision, user agency holds up. **Only 11% of consumers are willing to let AI decide their purchase, even in low-risk categories.** In the earlier stage, delegation is advancing fast. **56% feel comfortable delegating all communication with a brand to their assistant,** and around a third has already instructed it to prioritize some brands over others. That delegation rests on the assistant's supposed neutrality. That neutrality is not real (see point 06).

04

Emotional and symbolic dimension

The available evidence doesn't confirm a weakened emotional bond — it qualifies it. A positive interaction with a brand's chatbot significantly reinforces attachment to, and advocacy for, that brand. The assistant's anthropomorphism also follows a curvilinear pattern: **a moderately human identity generates more trust, an excessively human identity generates discomfort.** With one relevant caveat: almost all of this evidence studies the brand's own chatbot, not the brand-agnostic, general-purpose agent analyzed here. **What happens to the emotional bond when the interlocutor doesn't belong to the brand remains unstudied.**

05

Temporality of relationships

Conditional loyalty already has numbers behind it. **A quarter of consumers plan to set their brand preferences through AI on a regular basis within a year, and 16% would already act on their assistant's recommendation without further deliberation.** The risk isn't just a more volatile relationship. If the brand doesn't make its value transparent to that system, loyalty can shift from the brand to the AI assistant itself.

06

Ethical and systemic scalability

The agent's neutrality has a documented problem. A Princeton study found that AI shopping models frequently recommend sponsored options, to the point that the US FTC warns this could constitute a deceptive practice. **A Harris Poll survey found that 75% of consumers would trust an AI agent less if they knew its recommendations were sponsored.** The market for infrastructure providers already shows its first major players. The question isn't whether a de facto oligopoly will exist — it's how one that is already starting to form gets regulated.

Market Trends

The upward market trend against the absence of a trust-architecture model

Size and investment

The agentic AI market goes from \$7.29 billion in 2025 to \$9.14 billion in 2026, at an annual growth rate of 40.5%. A large part of that capital isn't aimed at making brands' back office more efficient. It's aimed at building the layer that mediates the relationship with the user. **McKinsey projects that agentic commerce alone will move a trillion dollars in the United States by 2030, and there are already more than 90 startups building on the three protocols — from Google, Anthropic, and OpenAI/Stripe — that standardize how an agent communicates with a brand.**

The gap

That growth coexists with a wide gap between what gets piloted and what reaches production. Gartner calculates that more than 40% of agentic AI projects will be cancelled before 2027, and that only 23% have scaled successfully. MIT's "GenAI Divide" report finds that 95% of generative AI pilots fail to deliver the expected return.

The root of the problem


That gap isn't explained by limitations of the model. A recent analysis of agentic commerce puts it this way: **the decision isn't which AI to buy, it's how the trust architecture gets built.** The model is the cheap part of that architecture. Another study on agent governance finds that failure almost always occurs outside the model: **undefined autonomy thresholds, undesigned escalation points to a human, absent monitoring.** It's an unresolved design gap, not a technical capability gap.

Sector Benchmark

Sector Benchmark

The mobility sector already shows the emerging model working — an agent that answers to the user, not the brand. Banking shows the opposite, even though it may look otherwise at first glance: its AI infrastructure is just as advanced, but the agent still belongs to the bank. Health and energy show that, without enough regulatory or transactional pressure, the same technology stays inside the institution and still doesn't mediate the relationship with the end user.

Sector	¿What happens today?	¿Who controls the agent?	Maturity	Key Data Point
Mobility	The agent searches, books, and pays for a complete trip (flight and hotel) within a single conversation, without the user visiting the airline's or hotel's website separately. There are also already agents that negotiate directly with the dealership when buying a car, on the buyer's behalf.	The user	Advanced	Sabre–PayPal–MindTrip alliance for travel bookings; car-buying agents with an average saving of \$1,500
Banking	The bank has built its own assistant to handle customer matters (queries, account management, some advisory).	The bank, not the user	Advanced	Bank of America's Erica has accumulated 3.2 billion interactions
Health	AI already manages clinical documentation, records, and authorizations within hospitals and insurers.	The health system, not the patient	Early-stage, concentrated in internal processes	Only 3% of health systems use AI in processes that directly touch the patient.
Energy	AI manages the power grid and building energy consumption, on the utility's side.	The utility, not the user	The least advanced of the four for this thesis	The idea of an agent that switches tariffs on the user's behalf has been researched for more than a decade, but has never become a real product.



Mobility

Here, the purchasing agent works for the user, not to the airline or the hotel. In February 2026, Sabre, PayPal, and MindTrip connected their systems so a user can search, book, and pay for a complete flight and hotel within a single conversation, with real-time inventory across more than 420 airlines and two million properties. IDC projects that, by 2030, 30% of travel bookings will no longer just be suggested by an agent, but completed by it. In car buying, the same logic already produces measurable results: services like CarEdge contact dealerships directly and negotiate the price on the buyer's behalf, with an average saving of \$1,500. In all three cases, it's the user who instructs the agent and benefits from its work.



Banking

Here, the investment is just as intense, but the agent still belongs to the bank. Erica, Bank of America's assistant, has now accumulated more than 3.2 billion interactions, and more than half of financial institutions already pilot or deploy agentic AI at an advanced stage. But that assistant resolves queries and account management within the rules the bank itself defines. It doesn't compare a competitor's offer or negotiate on the customer's behalf against another bank. It's the current model, faster and more efficient, not the emerging model: the power to decide what's offered to the customer still sits with the brand — just automated now.

Start free trial

Contact Sales

■ UNMATCHED CAPABILITIES

Complete capabilities.
Total control.

Fin's industry-leading capabilities let you analyze, train, test, and deploy, all in one workspace. It's designed to help you fine-tune quality and drive continuous performance improvements—no code required.

View demo

Start free trial

Guidance

Learn

Basics

Tone of voice

Professional

Answer length

Standard

Context and clarification

Create unique guidance on the follow-up questions Fin should ask.

Check device type

Enabled

Used: 2131 • Resolved: 86% • Routed: 14%

If a customer reports a bug, first clarify whether they are using the iOS or Android app before proceeding with troubleshooting and support. Never provide an answer without having this information.

Optimize

Audience

Cancel

Pause

Save

Analyze

Monitor, analyze, and optimize Fin by spotting trends, filling content gaps, and fixing quality issues topic by topic.

Train

Customize Fin's tone of voice, teach it your support policies, and configure how it handles complex tasks in over 45 languages.

Test

Test answers, review the sources and settings that shape them, and get tailored recommendations to optimize performance.

Deploy

From email and live chat to phone, SMS and social - Fin can answer any question, across any channel, any time.

Health

Here, the agent that decides works for the hospital or the insurer, not the patient. 75% of health systems have already deployed some AI solution, up from 59% a year ago. **But that adoption is concentrated in eligibility verification, denial prediction, and clinical decision support — internal processes of the health system. Only 3% have brought it into real workflows;** the rest remains in pilot phase. The push toward an agent that directly represents the patient in front of their insurer exists, with regulatory programs such as the \$50 billion rural health program pursuing exactly that, but it's still an institutional announcement, not a deployed reality. A notable example would be Hippocratic AI.

Luz + gas

Solar

Movilidad Eléctrica

Aerotermia

Empresas

Ayuda

Mi cuenta

KRAKENTECH

Kraken: Nuestra arma secreta

Kraken es nuestra plataforma tecnológica en la nube, una herramienta capaz de conectar todos los pasos de la cadena de abastecimiento, desde la generación de energía renovable, hasta la atención al cliente y la facturación. Las mentes del equipo de Octopus la han diseñado y desarrollado exclusivamente para gestionar todos los aspectos del suministro eléctrico de nuestros clientes. Y la siguen modificando para facilitar al máximo sus vidas.

Octopus nació con el objetivo de cambiar el sector energético, por eso teníamos que hacer las cosas de manera diferente. Y lo hicimos gracias a Kraken, convirtiéndonos en una empresa tecnológica, en la que todos los procesos están automatizados e integrados. Es el resultado de un trabajo en equipo que ha permitido a Octopus convertirse en la empresa líder en el sector energético.

Energy

Here, the agent that manages the grid works for the utility, not the user. The agentic AI market in energy grows from \$656.6 million in 2025 to a projected \$14.9 billion in 2035, but that investment is directed at grid self-healing, demand response, and energy trading between facilities — on the company's side. **The precedent conceptually closest to this document's thesis, an agent that switches tariffs on the user's behalf, is also the one that has advanced the least: it has been researched academically for more than a decade, and has never become a real product.** A notable example would be Kraken Tech, from Octopus Energy.

Possible Scenarios

Three possible scenarios for addressing the trust gap

Scenario 1

The API-fication of the brand

The intent already exists: route the purchase process to a generic agent to speed it up, without the brand having to build its own conversation with the user. OpenAI opened that path in ChatGPT, with Etsy connected from the start and a million Shopify merchants announced, including Glossier and Skims. Months later, only a dozen had actually connected: what's missing is the technical groundwork — real-time inventory, price, and payment — not the will to do it. It's a scenario that, if resolved, would change the relationship between brand and user at the root: the purchase would stop happening in the brand's channel and start happening inside the agent.

Scenario 2

The war of the agents

Explicit rules exist between agents. It's no longer a hypothesis: in April 2026, Anthropic let Claude agents negotiate 186 deals across more than 500 items, transacting just over \$4,000 in real money with no human intervention. Services like CarEdge already deploy an agent that negotiates with dealerships on the buyer's behalf, with an average saving of \$1,500. High-consideration purchases — cars, homes, travel — are the categories in which this scenario is likely to emerge first.

Scenario 3

The renaissance of the brand

The brand as an anchor of human trust, not as an immersive experience. This is no longer a future projection, as the delegation data in our Hypothesis already showed. The mechanism is simple: faced with a system the user doesn't fully control, the familiar brand reduces the risk of delegating. It is, literally, the example of "a specific brand's sustainable sneakers" that the original document raised, only now it's already happening.



The future points to an intermittent system among the three scenarios,
where the brand acts as a trust filter over an infrastructure and a
negotiation that the user cannot fully oversee

Design Principles

four design principles

What has been described so far does not call for a better interface. **It calls for revisiting the discipline of product design on four fronts:** who we design for, what we design, how we build trust, and what part of the relationship with the user we still control.

Design for two audiences: the person, and the agent that represents them

Until now, a product's audience was a person. Today it splits: part of the decisions now first pass through an agent acting on that person's behalf.


Designing for the agent...

...requires structured, complete, unambiguous data: price, stock, return conditions, in a format the agent can read without interpreting. But that agent is not a neutral reader. Designing for the agent is no longer just about being clear for a machine. It's knowing that machine isn't impartial, and that how data is presented can weigh as much as the data itself.

Designing for the human

...changes purpose. The interface that used to serve discovery and exploration now serves to validate and oversee what the agent has already decided on the user's behalf.

The product designer no longer writes only for the person who decides. They also write for the system that can tilt that decision, depending on how the data is worded.



A brand's value proposition must be legible and quantifiable

It is no longer enough for a catalog to be well photographed and well written.

Data and brand signals...

...have to become legible to a system: price, stock, and conditions, in a format that requires no human interpretation. This also extends to attributes that used to live only in the brand's emotional territory. **"Sustainability," "craftsmanship," or "good service" have to translate into third-party certifications, standardized scores, or verifiable metadata.** Whatever can't be encoded that way, for an agent, simply doesn't exist.

The gap between intention and execution...

...shows up here too. Declaring that a brand wants to be legible to these systems is easy. **Keeping that data up to date, free of contradictions across channels and free of ambiguity,** is ongoing engineering and content work, not a box that gets checked once and forgotten.

Designing the value proposition isn't just a content-strategy exercise. It's a product-architecture decision, one that demands technical and operational groundwork



Design the trust architecture, not just the interface

The market evidence already says it: the model isn't the problem. Clear autonomy thresholds are missing. It's still undefined when an agent should stop and escalate the decision to a person. That trust architecture is product design's job, not a decision that happens in another department.

Auditability...

...turns a black box into a reviewable decision. **The interface must be able to answer, when the user asks, a specific question: why did my agent choose this option over the other one?** Showing the result isn't enough — the criteria have to be shown too.

Calibrated friction...

...contradicts the usual instinct to eliminate all friction. **Before a high-stakes decision — a large purchase, a negotiation closed on the user's behalf — introducing an explicit confirmation step, or a moment of doubt before the decision is finalized, is not an obstacle.** It's what returns control to the user they would otherwise have delegated without knowing it.

Neither piece exists by default today in most agentic systems. Designing them is, according to the market evidence itself, the difference between a pilot that gets cancelled and a system that reaches production.



Reclaim authorship of the message **within the conversational channel**

With social media, the brand lost control of the sales channel, but not of its own message: the photo, the video, the tone, remained its own. In the conversational channel, the channel and the message become separated for the first time. When an agent summarizes a product, the user doesn't see the brand's campaign — they see a text written by the model itself, with its own judgment about what to highlight. The brand doesn't lose the channel. **It loses authorship of what's said about it.**

Reserving your own space within someone else's channel...

When a brand builds its own application within the conversational environment, instead of letting the model write the summary, it reclaims part of that authorship. **It's a product design decision, not a commercial negotiation.**

Designing the material that feeds that summary...

If the agent is going to write on the brand's behalf, **the brand can at least decide what words, emphasis and tone inform the summary**, instead of leaving it to the model's own judgment.

Existing research studies the brand's own chatbot. What's missing is an answer to what happens when that assistant doesn't belong to the brand, but to a third party.



Takeaways

01

Delegation has a hidden edge

Users delegate research and comparison, but the final call stays human, and that decision is more exploitable than it looks: an agent's ranking can be nudged by how a listing is written, and no rule today requires that nudge to be disclosed to the user.

04

Most "agent-readiness" work is translation, not engineering

Turning "sustainable" or "well-made" into a certification or a metadata field a machine can check is mostly a content and governance problem. Skipping that step is a likelier reason a pilot gets cancelled than any limitation in the model itself.

02

The gap between "let it filter" and "let it decide"

More than half of users already let an agent filter what they see from a brand, but only 11% would let it close a purchase outright. That gap between filtering and deciding is where most of the design work in this paper lives.

05

Sector speed tracks bargaining power, not tech investment

The sectors furthest along didn't get there through AI investment: banking's AI infrastructure is just as advanced as mobility's. They got there because users already had aggregators and easy switching before agents existed. That's a better predictor of who mediates the relationship next than any AI budget line.

03




The biggest unknown isn't technical, it's emotional

A good chatbot interaction still builds attachment to a brand, but almost none of that evidence covers a brand-agnostic assistant like the one this paper is about. Whether that warmth transfers when the assistant belongs to nobody in particular remains the single biggest open question here.

**Mormedi can help you lead the transition
to intelligent products through strategy,
design, and trust-architecture innovation**

Our unique value

We help companies to innovate throughout all phases and product dimensions



Discover
Opportunity mapping

Create and Deliver
Value proposition and solution in detail

Implement
Go to market

Clarity in complexity

We combine business vision and rigorous technological understanding with a "Customer centric" approach to formulate relevant and attractive value propositions.

"Mormedi translated complex systems into simple, effective narratives and value propositions which are easy to understand, even for a non-technical audience."

Creativity with Purpose

Mormedi crafts seamless, consistent experiences across physical and digital touchpoints, ensuring that customers can interact with your products and services when, where, and how they choose.

"Mormedi are a highly creative team that deeply understood our needs and turned them into value with a systemic approach."

Real world implementation

We have helped leading brands all over the world to implement their solutions, from engineering support to creating powerful commercial narratives and "go to market" strategies.

"Mormedi's diverse skill set provided critical expertise from strategy all the way to implementation."

End-to-end delivery

From insight-identification and concept exploration, to develop, build, and launch, we have a team that can deliver value at every stage of the process, ensuring what is launched to the market is what you need to achieve your commercial goals.

Proven methodology

Agilely delivering and developing designs through an iterative process, that starts with building confidence and engagement within your organization and ends with the definition of world-class product experiences.

Diverse team


With a global perspective and varied expertise supplemented by a network of sector and region experts. We take active steps to ensure our work is inclusive, representative, ethical, and sustainable, as well as commercially viable.

Experience in delivery

Our multi-sector approach allows us to bring the best practices to bear on every project we undertake. That, allied to a process underpinned by consumer research and validation means you can be confident in the validity of our work.

We have worked for 22 Fortune 500 companies

We are trusted by both industry giants and emerging leaders.
We build long-term partnerships grounded in shared impact.



Sources

Context

- Braze, Inc. (2026). 2026 Global Customer Engagement Review. Comunicado de prensa, 24 de febrero de 2026. <https://www.businesswire.com/news/home/20260224711754/en>
- Locus (2026). Q2 2026 US Consumer Survey ("New Locus Survey Finds Shifting Shopper Preferences Are Creating New Retail Logistics Pressures"). PRNewswire, 3 de junio de 2026. <https://www.prnewswire.com/news-releases/new-locus-survey-finds-shifting-shopper-preferences-are-creating-new-retail-logistics-pressures-302789418.html>
- Anthropic (2024). Introducing the Model Context Protocol (MCP). 25 de noviembre de 2024. modelcontextprotocol.io
- Google Cloud (2025). "Announcing Agent Payments Protocol (AP2)." Google Cloud Blog, 16 de septiembre de 2025 (menciona A2A y MCP como protocolos de base sobre los que se construye AP2). <https://cloud.google.com/blog/products/ai-machine-learning/announcing-agents-to-payments-ap2-protocol>
- OpenAI (2025). "Buy it in ChatGPT: Instant Checkout and the Agentic Commerce Protocol." 29 de septiembre de 2025 (ACP, co-desarrollado con Stripe). <https://openai.com/index/buy-it-in-chatgpt/>
- CiteLens (2026). "CiteLens Study: SEO Decides AI Citations on Google and Perplexity, Not ChatGPT." 8 de julio de 2026 (cifras exactas: 93% Google AI Mode, 89% Perplexity, 30% ChatGPT, sobre 320 consultas). <https://www.einpresswire.com/article/925230382/citelens-study-seo-decides-ai-citations-on-google-and-perplexity-not-chatgpt>

Possible Scenarios

- OpenAI (2025). "Buy it in ChatGPT: Instant Checkout and the Agentic Commerce Protocol." 29 de septiembre de 2025. <https://openai.com/index/buy-it-in-chatgpt/>
- CNBC (2026). "OpenAI revamps shopping experience in ChatGPT after struggling with Instant Checkout offering." 24 de marzo de 2026 (confirma el repliegue y el bajo número de comercios realmente conectados). <https://www.cnn.com/2026/03/24/openai-revamps-shopping-experience-in-chatgpt-after-instant-checkout.html>
- Anthropic (2026). "Project Deal: our Claude-run marketplace experiment." 24 de abril de 2026 (experimento realizado en diciembre de 2025). <https://www.anthropic.com/features/project-deal>
- CarEdge (2025-2026). Datos de la propia compañía, recogidos por 10News San Diego, "AI-powered car buying service helps consumers save money and avoid dealership stress", 6 de octubre de 2025. <https://www.10news.com/news/local-news/ai-powered-car-buying-service-helps-consumers-save-money-and-avoid-dealership-stress>(nota: la propia web de CarEdge actualiza esta cifra con frecuencia, en julio de 2026 hablaba de un ahorro medio de \$2.487; conviene confirmar qué cifra exacta se quiere citar y con qué fecha, dado que es un dato que la compañía refresca).
- Ver Gale (2026), The Preference Economy, ya referenciado en el punto 01 (Hypothesis). Es la misma fuente.

Hypothesis

- Allouah, A., Besbes, O., Figueroa, J. D., Kanoria, Y., & Kumar, A. (2025). "What Is Your AI Agent Buying? Evaluation, Biases, Model Dependence, & Emerging Implications for Agentic E-Commerce." arXiv:2508.02630 (versión revisada dic. 2025; aceptado en Proceedings of the ACM Web Conference 2026). <https://arxiv.org/abs/2508.02630>
- Gartner, Inc. (2026). "Gartner Survey Finds Consumers Want AI Shopping Help, But Not AI Purchase Decisions." Comunicado de prensa, 27 de mayo de 2026 (encuesta a 322 consumidores de EE. UU., enero de 2026). <https://www.gartner.com/en/newsroom/press-releases/2026-05-27-gartner-survey-finds-consumers-want-ai-shopping-help-but-not-ai-purchase-decisions>
- Gale (2026). The Preference Economy: Earn True Loyalty or Get Filtered Out. Encuesta a 3.000 consumidores de EE. UU. y Reino Unido. Cobertura primaria: Retail Dive, "Why marketers must rethink loyalty as AI reshapes consumer connections", 4 de junio de 2026. <https://www.retaildive.com/news/marketers-rethink-loyalty-ai-reshapes-consumer-connections/822305/>

Design Principles

- "Same position-effects study" (Principle 1) → Allouah et al. (2025), ver sección 01.
- "Gartner, project cancellation" (Principle 3) → Gartner (2025), ver sección 02.
- "Trust-architecture analysis" (Principle 3) → sin fuente primaria identificada todavía; el texto del paper lo describe como "a recent analysis of agentic commerce" sin nombrar organización. Si tenéis el artículo original a mano, dadme el título o el enlace y lo formateo.
- "Agent governance study" (Principle 3) → mismo caso: necesito el nombre del estudio o una pista (autor, plataforma) para localizarlo y verificarlo antes de darle formato de cita completa. No quiero inventar una atribución.

Market Trends

- Fortune Business Insights (2026). "Agentic AI Market Size, Share & Forecast Report, 2034." <https://www.fortunebusinessinsights.com/agentic-ai-market-114233>
- Gartner, Inc. (2025). "Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027." Comunicado de prensa, 25 de junio de 2025. <https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027>(nota: la cifra del 23% que "ha escalado con éxito" circula ampliamente atribuida a McKinsey en análisis de mercado de 2026, pero no he podido confirmar el informe primario exacto de McKinsey que la contiene; si la usáis, recomiendo verificarla en el informe de McKinsey citado más abajo o pedir la cita directa al analista que la incorporó).
- Challapally, A., Pease, C., Raskar, R., & Chari, P. (2025). The GenAI Divide: State of AI in Business 2025. MIT NANDA (Project NANDA, MIT Media Lab), julio de 2025. https://nanda.media.mit.edu/ai_report_2025.pdf
- ICSC & McKinsey & Company (2026). "Shopping in the Age of AI: Redefining Stores for a New Era." 27 de abril de 2026 (encuesta a 3.004 consumidores de EE. UU., enero de 2026). <https://www.mckinsey.com/industries/retail/our-insights/shopping-in-the-age-of-ai-redefining-stores-for-a-new-era>

Sector Benchmark

- Sabre Corporation, PayPal Holdings, & Mindtrip (2026). "Sabre, PayPal, and Mindtrip Partner To Deliver the Industry's First End-to-End Agentic AI Experience for Travel." Comunicado de prensa conjunto, 12 de febrero de 2026. <https://www.prnewswire.com/news-releases/sabre-paypal-and-mindtrip-partner-to-deliver-the-industrys-first-end-to-end-agentic-ai-experience-for-travel-302686303.html>
- Bank of America (2026). "BoFA AI and Digital Innovations Fuel 30 Billion Client Interactions." Comunicado de prensa, 10 de marzo de 2026. <https://newsroom.bankofamerica.com/content/newsroom/press-releases/2026/03/bofa-ai-and-digital-innovations-fuel-30-billion-client-interacti.html>

Contact:

Álex Mena
Lead Digital Product Designer
alexandra@mormedi.com

Thank you ;)