BUSINESS-LED. PEER-POWERED COMMUNITY.

Fin.Tech MARKETING

What Came Out of Singapore: AI Ethics, Brand Authenticity, and the End of Surface-Level Personalization

Fin.Tech Marketing Community Roundtable, 5 March 2026

The Fin.Tech Marketing Community met in Singapore on 5 March to talk about AI. Specifically: what's actually working, what's still hype, and where the ethical lines are when marketing gets personal.

The big takeaway? AI is a tool for getting better at what you already do. It is not a strategy on its own. The group kept coming back to the same point: figure out the business problem first (productivity, revenue, risk), then decide whether AI helps solve it.

The Room's Agenda

The conversation covered three areas that kept overlapping: how personalization needs to go deeper than first-name tokens, why transparency about AI use is becoming a competitive requirement, and how search is shifting now that LLMs like Gemini and ChatGPT are recommending brands based on what the rest of the internet says about them.

On personalization, the group pushed back on the "hyper-personalization" label. Real personalization means contextual relevance based on what a customer is actually doing, not just who they are. And in fintech, where trust is everything, that means disclosing when AI is involved and keeping data out of public models.

On retention vs. acquisition: several participants made the case that personalization for existing customers, based on their specific usage patterns and needs, matters more than top-of-funnel flash. That's where you prevent churn and build long-term relationships.

And on the generative vs. agentic distinction: most teams are winning with generative AI right now (faster content, fewer bottlenecks). The next frontier is agentic AI, where tools don't just draft the email but decide when to send it and why.

From the Speakers

Robin Lee kept bringing the conversation back to ethics. His main point: the "creepy factor" in personalization is real, and marketers need to find the line between helpful and intrusive. He shared a use case worth noting: building an AI agent trained on specific brand guidelines to act as a final quality check on tone and voice. He also warned about "generic sameness." If every team uses the same prompts, every brand starts sounding identical. His solution: creative prompting, like writing in the voice of a specific character, to keep brand voice distinctive. Robin's upcoming book, The Full Stack AI Ethicist, goes deeper into guardrails for teams handling sensitive financial data.

Moira Ann focused on the distinction between generative and agentic AI. Her argument: most teams are comfortable with AI-assisted content creation, but the real shift is toward AI that can pursue complex goals on its own. She also flagged that your website isn't the only place prospects evaluate you anymore. AI models pull from across the web to form recommendations, which means your external signals (press, blogs, community presence) matter as much as your homepage. On governance, she stayed firm: AI can flag a risk, but a human needs to make the regulatory call.

What's Next

If you're working on any of these problems (brand governance for AI-generated content, deeper personalization, or building internal AI agents), share what you've learned with the community. The best ideas at this roundtable came from people who tried something, failed, and told us about it.

Always Fin.Tech.