MASTER'S THESIS · FINAL DEFENCE

Generative AI in Product Discovery

Adoption patterns · Perceived value · Barriers · Role evolution among Product Managers

CANDIDATE

Yash Sudhir Todkari

SUPERVISORS

Prof. Marco Cucculelli (UNIVPM) · Prof. Olgun Aydın (Gdańsk Tech)

Università Politecnica delle Marche · Masters in International Economics & Commerce · A.Y. 2025–2026

THE GAP & THE QUESTIONS

A blind spot where two mature fields meet

GenAI-in-knowledge-work and product management are each well studied — but no academic work examines how PMs adopt GenAI specifically **within product discovery**. McKinsey projected routine organisational GenAI use rising ~33% (2023) → ~79% (2025), yet PMs adopt with little top-down coordination.

THE MAIN QUESTION

How are PMs integrating GenAI into product discovery — and what are the perceived impacts on efficiency, output quality and business outcomes?

1 SRQ1
Adoption

Which tools & use-cases — at what maturity?

2 SRQ2
Perceived value

What value does GenAI bring to discovery?

3 SRQ3
Barriers & trust

What limits effective, trusted use?

4 SRQ4
Role evolution

How are skills & the PM role changing?

FRAMEWORK & METHOD

Three lenses, a convergent design, honest limits

1 TAM / UTAUT

Why adoption happens — usefulness, ease, facilitating conditions.

2 Augmentation theory

Daugherty & Wilson — AI redistributes cognitive labour.

3 AI Fluency (4D) ★

Delegation · Description · Discernment · Diligence — first academic application.

CONVERGENT PARALLEL MIXED-METHODS

82 practitioner survey

PMs & Product Owners · closed items → descriptive stats & cross-tabs · open items → thematic analysis

4 firm case studies

Figma · Notion · Miro · Atlassian — organisational-level evidence, triangulated against the survey.

SCOPE & LIMITS

Exploratory & convenience-sampled (LinkedIn / PM communities); read as analytical, not statistical, generalisation. Cross-tabs are descriptive — no significance testing; Likert means are directional summaries; small sub-group cells (e.g. scale-up n = 12) are indicative only.

CASE STUDIES · ORGANISATIONAL LENS

Four leading firms — five consistent themes

	Figma	Notion	Miro	Atlassian
01	**Research synthesis & documentation**	The #1, most measurable value driver across all four firms (Notion: 7+ hrs/week/team).		
02	**Augment, don't replace**	Every firm positions GenAI as augmenting human judgement — in product and guidance.		
03	**Transparent limits → discernment**	Firms openly disclose bias, convergence risk & context gaps; outputs are starting points.		
04	**Governance as an enabler**	Encryption, data residency, BYOK, tiered access — conditions that lower enterprise barriers.		
05	**Discovery→delivery toolchains**	MCP integrations let discovery context flow into engineering IDEs (Cursor, Copilot, Claude Code).		

FINDING 1 · SRQ1 — ADOPTION

Routine, broad, and led by general-purpose LLMs

82
PMs & POs surveyed

93%
Daily / Several × week

56
Use it daily

10
Activities in active use

Most-used GenAI tools

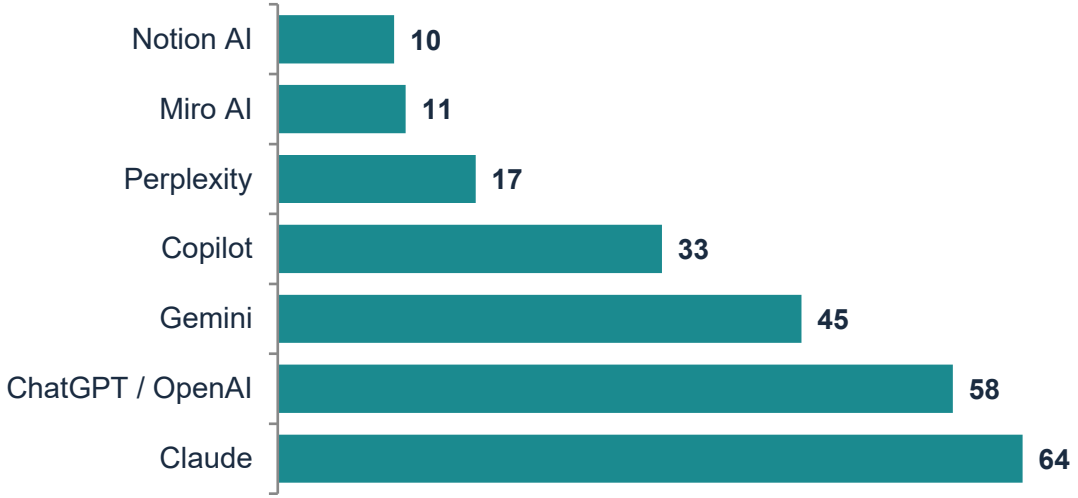

Top discovery activities supported

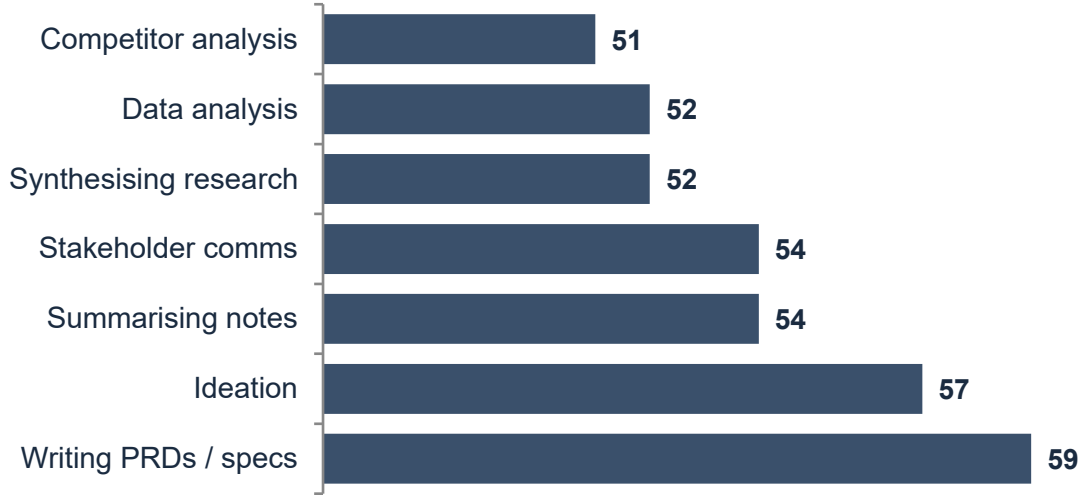

Sample skews senior (48% Senior PM; 81/82 with 1+ yr) and enterprise — engaged adopters more than the full PM population.

FINDING 2 · SRQ2 — PERCEIVED VALUE

Useful, yes — but trusted less

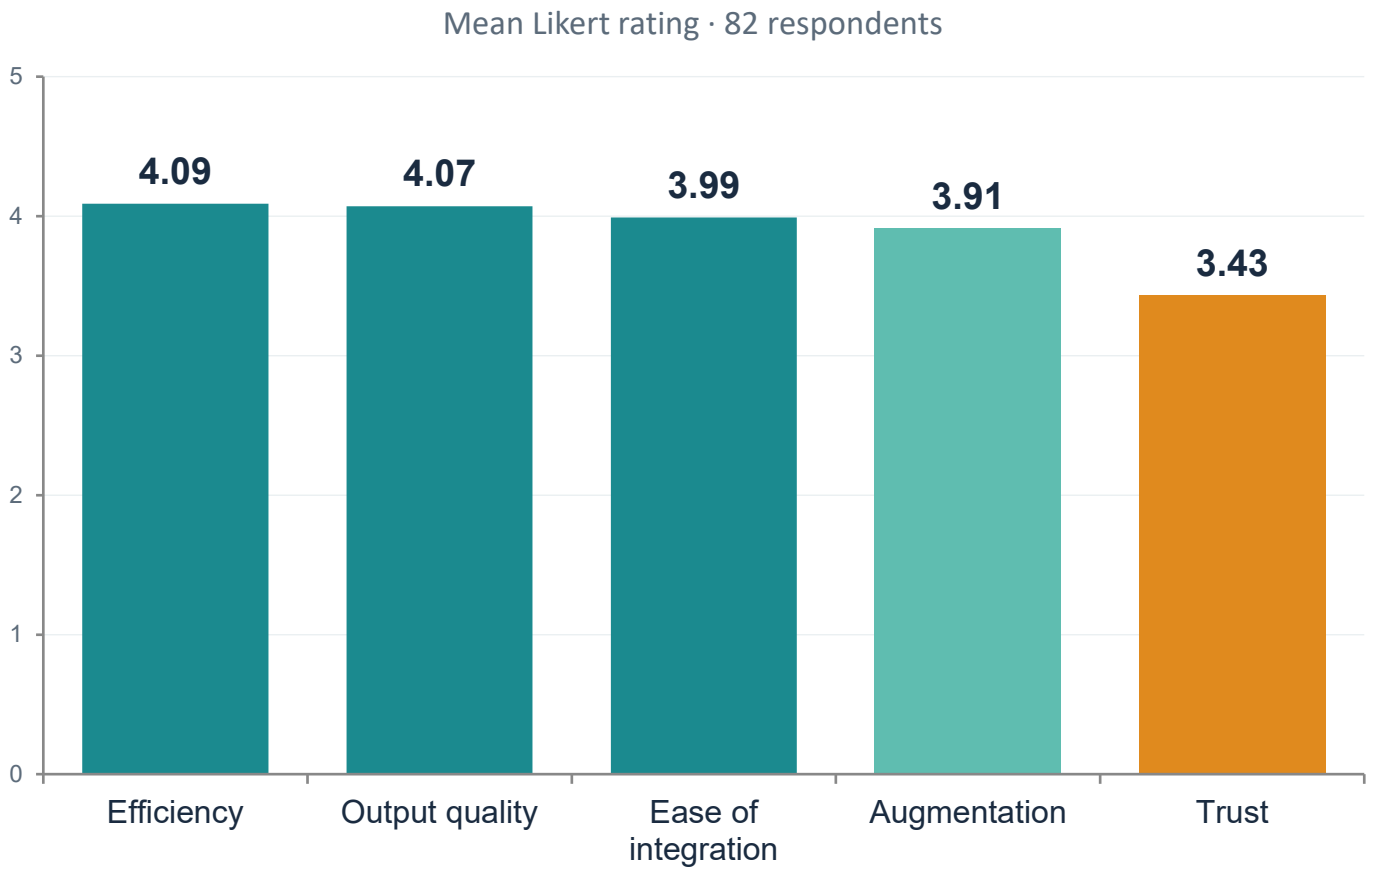

THE PATTERN

Usefulness items cluster around 4.0; trust is lowest at 3.43. In agreement terms: **78% agree it is efficient — only 48% agree they trust its output.**

WHAT IT MEANS

Adoption is pragmatic and conditional. PMs value what GenAI lets them do — but still own and verify the output. It is a fast productivity & synthesis tool, not an autonomous decision-maker.

FINDING 3 · SRQ3 — BARRIERS & TRUST

Barriers are about reliability & governance — not usability

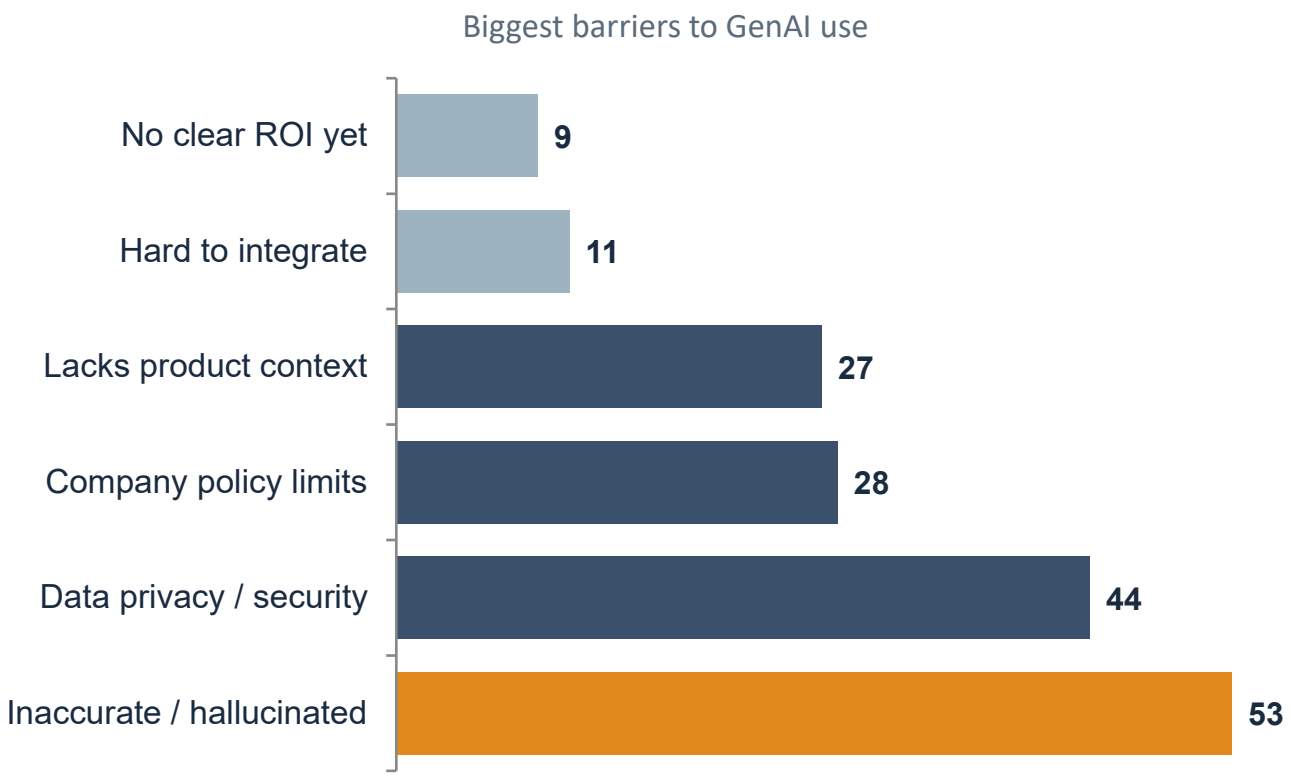

TRUST IS CONDITIONAL, NOT ABSENT

Extended selectively — by task stakes, available context and governance. Human review is a structural requirement, not a transitional inconvenience.

THE DEEPER BARRIER

Hallucination is named most — but the open text reveals **missing organisational context** as the harder problem. Better models alone can't fix it.

INDUSTRY CORROBORATION Deloitte (State of GenAI in the Enterprise, Q3 2024): 68% of firms have moved ≤30% of GenAI pilots into production; top barriers are regulatory, risk & governance-related — scaling and trust, not capability, are the bottleneck.

FINDINGS 4 & STRUCTURE

The gap is structural; the role is reshaped

THE TRUST GAP HOLDS ACROSS EVERY LENS

Lens	Usefulness	Trust
Role (5)	3.8–4.8	2.8–3.7
Company size (4)	3.8–4.4	3.0–3.5
Experience (5)	3.1–4.4	2.5–3.8
Region (4 major)	3.8–4.2	3.3–3.6

17/18 segments where trust is the lowest (or joint-lowest) of the five items.

The gap is a property of the practice, not of any one context (only exception: the 1–3-yr band). Read descriptively; cells indicative only.

SRQ4 · THE ROLE IS RESHAPED, NOT REMOVED

Manual Production → **Judgement & Orchestration**

Cognitive labour is redistributed — the model drafts and structures; the human judges, prioritises and stays accountable. The “strategic editor”.

SKILLS RISING IN VALUE (mentions)

CORROBORATION

Independent strands point the same way

PRACTITIONERS' OWN WORDS

> *"…turned days of synthesis into an hour-long task."*
>
> **Efficiency 4.09 · "time saved" top benefit**

> *"Generic excellence, but it lacks our internal context."*
>
> **Barriers: inaccuracy 53 · lacks context 27**

> *"From content creator to strategic editor… you can't outsource judgment."*
>
> **Trust 3.43 · augmentation 3.91**

Independently coded themes map almost one-to-one onto the closed-question results — when two strands agree, confidence rises.

ONE LOGIC · THREE LEVELS

MACRO **Literature**

GenAI strongest in synthesis; adoption governed by usefulness, facilitating conditions & augmentation.

MESO **Case studies**

Four firms operationalise exactly this — synthesis features, augmentation framing, governance.

MICRO **Survey**

Practitioners already work this way — synthesis-led use, usefulness above trust, human review insisted upon.

★ PRINCIPAL ORIGINAL CONTRIBUTION

The Refined GenAI Adoption Maturity Model

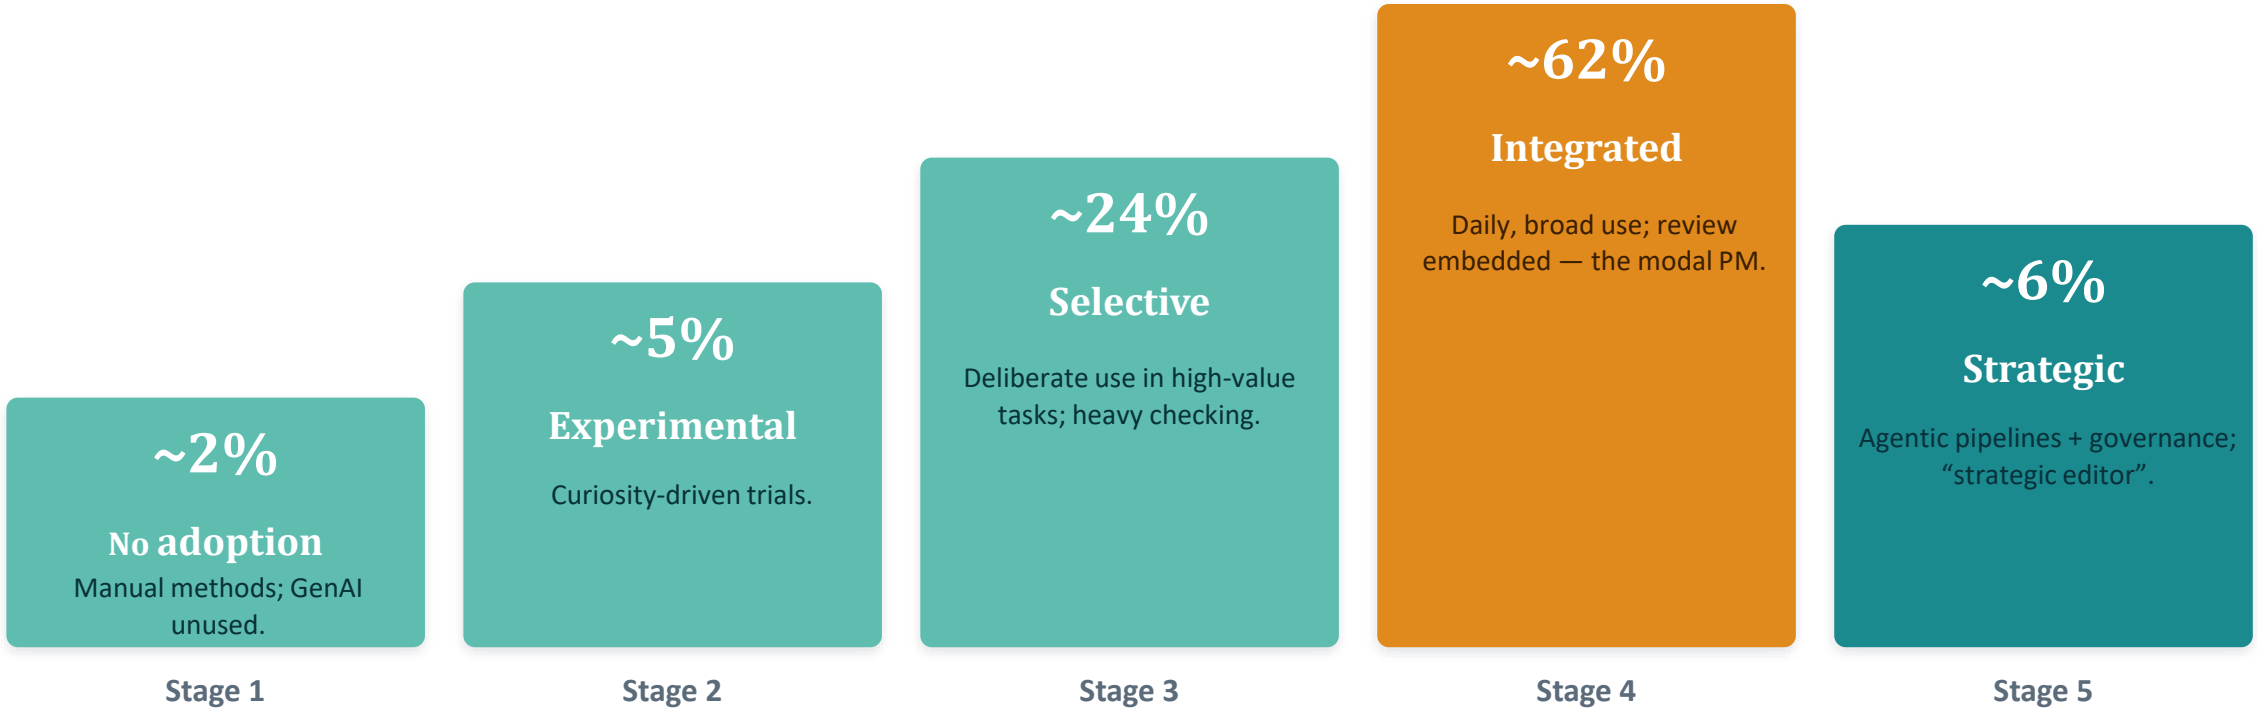

*"Teams advance not by using AI **more**, but by using it with greater **discernment, calibrated trust, and governance maturity**."*

— Yash Todkari · GenAI Adoption Maturity Model (2026)

THE BOTTOM LINE — BEYOND “AUGMENTATION”

Three findings that invert what the industry assumes

1

Trust plateaus by design

DATA Trust (3.43) stays below usefulness in every segment — even among veterans. The binding barrier is missing context, not model quality.

SO WHAT Design for permanent, calibrated human-in-the-loop.

2

The bottleneck is the scale-up

DATA Scale-ups post the weakest profile (trust 3.00); enterprises are steady (3.50). Friction sits in the rapid-growth middle.

SO WHAT Target enablement at scale-ups, not only enterprise compliance.

3

Maturity is discernment, not usage

DATA 62% already use GenAI daily and broadly. More usage no longer differentiates teams.

SO WHAT The moat is context-injection & calibrated trust — not tool access.

ADOPTION IS SETTLED · DISCERNMENT IS THE FRONTIER

*“The advantage no longer belongs to the teams that use AI the **most** — but to those who have learned, precisely, **when not to trust it.**”*

— The Discernment Principle · Yash Todkari (2026)

IN CLOSING

Adoption is settled.
Discernment is the frontier.

Usefulness is abundant; trustworthy judgement is the scarce resource. The firms that win build capability around automation — they do not let it stand in for judgement.

Thank you

I welcome your questions and feedback.

Yash Sudhir Todkari

Supervisors: Prof. Marco Cucculelli · Prof. Olgun Aydın · Masters in IEC, UNIVPM × Gdańsk Tech