



UNIVERSITÀ POLITECNICA DELLE MARCHE
FACOLTÀ DI ECONOMIA “GIORGIO FUÀ”

Bachelor's Degree Course in Economics (L-18, DM 270/04)

Final Report in Economics

On-Device AI Assistants and Consumer Purchasing Decisions

*A Behavioural–Economic Analysis of Smartphone Commerce, with
Implications for Markets and Firm Strategy*

Supervisor:

Prof. Matteo Picchio

Final Report of:

Elif Mustu S1118360

Academic Year 2025/2026

INDEX

Introduction	4
Historical evolution of mobile purchasing interfaces.....	5
Methodology and approach	6
CHAPTER I — THEORETICAL BACKGROUND AND CONCEPTUAL FRAMEWORK	7
I.1 Consumer purchasing as a multi-stage process.....	7
I.2 Search costs, bounded rationality, and choice overload.....	8
I.3 Digital intermediaries, platforms, and market power.....	9
I.4 AI-mediated commerce in the marketing and IS literature	10
I.5 On-device versus cloud AI: a conceptual distinction.....	11
I.6 Involvement, risk, and the moderating role of product category	13
I.7 Propositions for future testing	13
CHAPTER II — ON-DEVICE AI ASSISTANTS: CAPABILITIES, ADOPTION, AND ECOSYSTEMS	14
II.1 Technical enablers and functional scope	14
II.2 A taxonomy of purchasing-related assistant functions	16
II.3 Adoption determinants and user trust	17
II.4 Comparative ecosystem analysis: Apple, Google, and Samsung	18
II.5 Integration depth and market practice.....	20
II.6 The European and Italian retail context.....	20
CHAPTER III — IMPACT ON CONSUMER PURCHASING DECISION-MAKING	22
III.1 Need recognition and problem awareness	22
III.2 Information search and consideration-set formation.....	23
III.3 Evaluation of alternatives and heuristic substitution.....	25
III.4 Purchase execution and payment integration	26
III.5 Post-purchase behaviour, loyalty, and returns.....	27
III.6 Pulling the stage-level effects together.....	28
III.7 Welfare and inter-stage dynamics	29
CHAPTER IV — ECONOMIC AND STRATEGIC IMPLICATIONS.....	30
IV.1 Market structure and ecosystem lock-in.....	30
IV.2 Implications for retailers and consumer brands	31
IV.3 Pricing, transparency, and competitive dynamics	33
IV.4 Macroeconomic channels and aggregate consumption	34
IV.5 Regulatory and consumer-protection considerations	34
IV.6 What each actor should do	36
IV.7 Regional policy, SMEs, and the Italian Marche context.....	37
CHAPTER V — DISCUSSION, LIMITATIONS, AND RESEARCH AGENDA	38
Synthesis of findings	38
Limitations.....	39
Future research agenda	39
Conclusions	40
Managerial and policy recommendations (summary)	41
Closing remarks.....	42

References	42
Bibliography	47
Footnotes	48
Appendix A — Capability comparison matrix.....	49
Appendix B — Literature search protocol	50
Appendix C — Mini-survey: instrument and results	51

INTRODUCTION

For most people the phone is now the main place where shopping actually happens, from the first idle search to the moment they tap "pay". For more than a decade that meant browsers, shopping apps, and a great many sponsored results. Around 2023 something different started to show up: generative AI built straight into the phone's operating system, able to read a plain-language request, sum up what a product is like, and even begin the purchase without the user ever opening a normal shopping app. Apple Intelligence, Google's Gemini Nano, and Samsung's Galaxy AI are the names most people have run into, but the bigger change underneath them is quieter. The computation is moving off remote servers and onto the chip in your pocket.

This report is about one slice of that change: on-device assistants, by which I mean OS-level agents that run at least part of their work locally, or are at least built to keep the speed- and privacy-sensitive parts on the phone. The distinction matters more than it first looks. We have had cloud assistants since Siri in 2011 and Google Now in 2012, and for shopping they never really amounted to much: too slow, a bit creepy, and bolted on rather than built in. Local models tip that balance. They make it feel cheaper to hand a small decision, say "find me running shoes under €120 that fit wide feet", to something that already sits right next to your personal data. That is not just a nicer gadget. It moves where the cost of searching falls, who gets to frame your options, and who ends up with the power in the market.

The question I keep circling back to is easy to state. How do on-device assistants change the way people actually buy things, and what does that mean for competition, for the firms involved, and for whether shoppers come out ahead?¹ Answering it means working out where in the buying process the assistant really bites, how it works in plain

economic terms, how it shifts power between platforms and the shops that sell through them, and what firms and regulators ought to do in response. The timing is not accidental either. The EU's Digital Markets Act, in force since 2022 and applied from 2023, is aimed squarely at the gatekeeper operating systems these assistants ride on (European Commission, 2022).²

I have kept the scope deliberately tight. The focus is ordinary consumer shopping on smartphones in developed markets, mostly the EU and North America, and I set aside business procurement, non-phone devices, and desktop chatbots except where they make a useful comparison. I also refuse to treat every purchase as the same thing: reordering coffee pods, choosing a €600 pair of headphones, and signing up for a financial product behave nothing alike, and I keep them apart throughout. One honest caveat before anything else: there is still very little peer-reviewed evidence on on-device assistant shopping, and the products themselves change every few months, so this is mainly a conceptual study, built on the existing literature and on the companies' own documentation (Apple Inc., 2024; Google DeepMind, 2024; Samsung Electronics, 2024) rather than a heavily measured one, though I do put a few of its claims to a small exploratory survey of my own (Appendix C). What I can offer is a clear way of thinking about the problem and a handful of claims specific enough to be tested properly later. Where I could not find trustworthy market-share numbers, I describe how the mechanism works instead of inventing a figure.

Historical evolution of mobile purchasing interfaces

It is worth remembering that mobile commerce never changed everything at once. Each wave of it really only fixed one part of the buying process and left the rest more or less alone. The earliest phones, roughly 2007 to 2012, just shrank the desktop shop onto a

small screen; you still searched and compared in a browser, and the thing that actually killed sales was the checkout, all that typing on a tiny keyboard with no card saved.

Retailer apps and marketplaces came next, across most of the 2010s, and their real trick was the logged-in one-tap purchase paired with the push notification that reminded you of a need you had half-forgotten (Grewal et al., 2016). Then, around 2020, social and live commerce arrived, and an influencer or a thirty-second clip started doing the job that expert reviews used to do.

The current wave, from about 2023, is different in kind. It drops a generative agent right at the operating-system level and lets you delegate several steps at once, just by talking.

An app can only ever show you its own catalogue, and a normal web search forgets what you asked a moment ago; the assistant does neither, holding the thread of a conversation and reading signals off your phone without shipping all of it to a server.

Line the waves up and each one chipped away at a different cost. Browsers cut the cost of physically getting to a shop. Apps cut the cost of logging in and finding your way around. Social commerce cut the cost of trusting a brand you had never heard of. The assistant goes after the last one standing, the mental effort of turning a pile of information into an actual decision. That is also why its effect is so lopsided: where an earlier wave already cleared the bottleneck, like a subscription app handling repeat orders, there is little left to gain, but where comparing options is still genuinely painful, as with a complicated durable, there is a lot. The Italian figures fit that story, with phone-based buying climbing fastest among the under-45s while the gap in digital skills between regions stubbornly refuses to close (ISTAT, 2024). My own hunch is that the gains land on better-off urban users first.

Methodology and approach

A word on how I went about this, since the method decides what the claims can and cannot carry. The backbone is a structured conceptual review rather than a systematic one in the PRISMA sense. I searched Scopus, Web of Science and Google Scholar for work published between January 2018 and early 2026, combining terms for the technology, the buying behaviour and the platforms; the exact strings are in Appendix B. I kept peer-reviewed articles, official regulatory texts and first-party technical documentation, and set aside undated blog posts, vendor marketing, and purely engineering papers with no consumer or economic angle. Because peer-reviewed work specific to on-device assistant shopping barely exists yet, I lean on adjacent literatures on purpose, voice commerce, chatbots, recommender systems and platform economics, and I read the manufacturers' own material against the grain for promotional bias. On top of that sits one piece of primary evidence: a small, exploratory survey of 40 smartphone users (Appendix C), built so that each item maps onto a specific mechanism in the framework. I treat it as a reality-check, not a test. The sample is self-selected and skews young, the measure is stated willingness rather than observed behaviour, and the analysis is deliberately simple, means, agreement shares, and a light open-coding of the free-text answers. So I read the survey throughout as consistent or inconsistent with the framework rather than as confirming it, and I keep the conceptual argument and the survey evidence clearly apart.

CHAPTER I — THEORETICAL BACKGROUND AND CONCEPTUAL FRAMEWORK

I.1 CONSUMER PURCHASING AS A MULTI-STAGE PROCESS

A purchase is rarely just the moment money changes hands. The classic account, from Engel, Blackwell and Miniard (1995) and still the one most textbooks open with (Solomon, 2020), breaks it into five steps: you notice a need, you go looking for information, you weigh the options, you buy, and then you live with the result and decide whether you would do it again. Going digital did not get rid of those steps. It squeezed them together and shuffled the order, so that an app now lets you search and compare in the same breath, and a social feed can plant the need in the first place, before you even knew you had it (Appel et al., 2020).

What assistants add on top of all this is delegation. You can hand a whole step over to the agent: ask it to read the reviews for you, filter by some constraint, or just suggest a default you will probably go along with. That agent is an odd sort of middleman. It is not a passive catalogue you browse, and it is not a human advisor whose motives you can read; what it suggests is shaped by its training data, by deals struck with the phone maker, and now and then by placement logic you simply cannot see. I use the five-stage model as a frame for the rest of the report and keep asking the same thing at each step: who is actually doing the work now, the shopper or the agent, and which numbers move as a result, how long the search takes, how many options get a look, what price gets paid, how often the thing comes back. The whole time I assume people are boundedly rational rather than perfect optimisers, in the sense I unpack next. Real journeys blur these tidy stages anyway (Puntoni et al., 2021): you might spot the need in a shop, search on the shop floor, chew it over on the bus, and finally buy from the sofa that evening.

I.2 SEARCH COSTS, BOUNDED RATIONALITY, AND CHOICE OVERLOAD

Stigler (1961) put the basic idea on a formal footing a long time ago: information about prices and quality is not free, and a rational shopper keeps looking only until the expected gain from one more check drops below what that check costs.³ Assistants go straight at the cost side of that trade-off. A single spoken question can pull together specs, ratings and availability that used to take half a dozen separate searches. And once the answer comes back almost instantly, fast enough that you barely register the wait, the small mental cost of asking just one more thing more or less vanishes. The rule of thumb that an interface stops feeling like waiting once it answers within a few hundred milliseconds goes back to early human–computer research on perception and response time (Card et al., 1983), and the intuition carries over neatly to talking with an agent. Simon (1955) supplied the other half of the picture with satisficing: faced with too much, people do not hunt for the single best option, they grab the first one that is good enough.⁴ An assistant that hands you a neat shortlist stamped "best for you" is really satisficing turned into a product feature. Whether that helps you comes down to how honest the shortlist is. If it is genuinely surfacing the best matches, you come out ahead. If it is tilting toward sponsors, or just falling back on the brands you always buy, you can end up worse off than if you had kept looking (Thaler and Sunstein, 2008). Pulling the other way is choice overload. Iyengar and Lepper (2000) famously found that giving people too many options can put them off buying altogether and leave them less happy with whatever they did pick. An assistant that boils a thousand listings down to three might rescue someone from that paralysis, or it might bury the oddball product that would actually have suited them best. The theory honestly does not tell us which way it tips, so I come back to it stage by stage in Chapter III.

1.3 DIGITAL INTERMEDIARIES, PLATFORMS, AND MARKET POWER

Smartphone shopping does not really happen in an open market. It happens inside a handful of platform ecosystems (Rochet and Tirole, 2003). Apple's iOS and Google's Android run the operating system, decide how apps are distributed, own the payment rails, and now increasingly control the default assistant sitting on top of all of it. The standard theory of two-sided platforms says these companies will subsidise one side, usually us, to make their money from the other, the advertisers and developers. Put an assistant into that structure and it becomes hard to see it as a neutral search box. It sits at a chokepoint, and that chokepoint is starting to look a lot like the ones we already fight about in general search and the app stores.

This is exactly why the Digital Markets Act went to the trouble of naming some platforms "gatekeepers" and loading them with duties such as interoperability and fair access (European Commission, 2022). Article 6, the part about default settings and self-preferencing, is what bites here. The moment an assistant nudges you toward the platform's own payment service or its affiliated marketplace, you are back in familiar abuse-of-dominance territory. I save the detail of what this does to switching costs for Chapter IV; the point for now is just that the assistant inherits a position of power it never had to earn.

I.4 AI-MEDIATED COMMERCE IN THE MARKETING AND IS LITERATURE

None of this is starting from a blank slate. Researchers spent the better part of a decade studying chatbots and voice shopping before on-device generative AI arrived, and some of what they found carries straight over. Luo et al. (2019) showed that a chatbot can genuinely lift conversion, but that the experience collapses the moment a question wanders off the script, which is still a fair description of how assistants behave at the

edges today. McLean and Osei-Frimpong (2019) surveyed people using voice assistants at home and landed on the same tension I keep meeting: the main reason people use them is plain usefulness, and the main reason they hold back is the sense of being watched. That tension is the seed of an idea I rely on throughout the report, which is that delegation depends heavily on the category. People will let an assistant reorder detergent long before they let it pick a pension.

For adoption specifically, information-systems research tends to reach for UTAUT, which boils take-up down to how useful and how easy something feels, plus a bit of social pressure and whether the supporting conditions are there (Venkatesh et al., 2003).

For assistants I would add two things that model underweights: trust, and privacy. Awad and Krishnan (2006) gave the underlying problem a good name, the personalisation–privacy paradox, where people want recommendations shaped around them but resent the profiling that makes the shaping possible. On-device processing is sold as the way out of that paradox, and whether it really is one is the question I take up next.

There is one more literature I lean on heavily. Work on recommender systems has shown for years that algorithmic suggestions push buyers toward already-popular items and pile sales up at the top (Ricci et al., 2015). If assistants are basically recommenders with a conversation bolted on the front, and I think they mostly are, then the small merchant who fails to feed them clean data is at real risk of simply disappearing from view.

1.5 ON-DEVICE VERSUS CLOUD AI: A CONCEPTUAL DISTINCTION

It is worth being precise about what "on-device" means, because the word gets thrown around loosely. When I say on-device assistant I mean an agent built into the operating system that does at least some of its thinking on the phone itself, as opposed to a cloud

assistant, which is really just a microphone and a screen handing everything off to a server. Almost every real product lives somewhere in between, in a hybrid arrangement that keeps the fast or sensitive parts local and sends the harder reasoning out to the cloud. And when I talk about lock-in later on, I mean the whole bundle that makes leaving expensive: the operating system, the installed apps, the saved payment setup, and now the assistant's accumulated memory of you.

The only reason any of this can run locally is hardware. Phones now ship with small, compressed language models and dedicated chips, NPUs, built for exactly this kind of work, and Apple's recent A- and M-series silicon and Qualcomm's Snapdragon platforms are marketed openly on their generative-AI muscle (Apple Inc., 2024; Qualcomm Technologies, 2024). Doing the work on the chip cuts out the round trip to a server and, at least in theory, keeps your calendar, your messages, and your browsing from ever leaving the phone.

I say "in theory" on purpose. On-device is not the same thing as offline, and it is a long way from being private by default. A hybrid system still ships the difficult queries to the cloud, usually after trimming or summarising them first, and a commercial deal can still drop a sponsored result in at that cloud stage even when the first pass ran locally. So I treat on-device as a direction of travel with very uneven implementations, not as a promise that the user is in charge. That ambiguity has a real economic edge to it.

Someone who used to turn down cloud-based personalised offers because they felt exposed might happily accept the on-device version, which widens the market for assistant shopping, even while it makes life harder for a regulator trying to follow data that keeps bouncing between phone and cloud with no clear handoff.

I.6 INVOLVEMENT, RISK, AND THE MODERATING ROLE OF PRODUCT CATEGORY

Not every purchase is equally easy to hand over. Marketing calls this involvement, the gap between a low-stakes habitual buy and a high-stakes decision that carries real financial, social or physical risk (Solomon, 2020), and I doubt assistants will flatten that gap any time soon. Coffee capsules, printer ink, a refill of shampoo: low-risk, bought again and again, which is exactly the useful, low-privacy-worry profile that McLean and Osei-Frimpong (2019) link to voice-assistant use. Clothes and electronics are trickier, because they carry the risk of looking wrong or not being compatible, so people will often take the search help while refusing to let the assistant make the actual call. And regulated products come wrapped in disclosures that a quick summary may not legally satisfy at all. So any blanket statement needs qualifying. "Assistants raise conversion" is mostly true where involvement is low and the integration is deep. "Assistants replace comparison shopping" only really holds for mid-involvement goods, and even then only when the summary is accurate. At the high-involvement end it is still mostly guesswork.

I.7 PROPOSITIONS FOR FUTURE TESTING

Pulling the review together, a handful of concrete, falsifiable claims fall out, and the later chapters lean on them. The search-time saving should be biggest for categories whose attributes are easy to digitise, and it should grow as the query gets more specific and the processing gets faster. For low-involvement buyers the shortlist should shrink the set of options they consider, whereas for others multimodal search should do the reverse and widen it. Heuristic substitution should climb whenever a review summary leaves out the links back to the original sources. Wallet-integrated checkout should lift conversion on cheap items, while at the same time raising the odds of an order the buyer

never quite meant to place. Preference memory on the device should push switching costs up over time, even against a rival that is genuinely cheaper. And consumer surplus, I suspect, will not rise smoothly with accuracy: past a certain error rate, the cost of returns simply swamps whatever got saved on search. I test these only lightly here, against the small survey in Appendix C; the proper testing waits for the kind of data the agenda in Chapter V describes.

CHAPTER II — ON-DEVICE AI ASSISTANTS: CAPABILITIES, ADOPTION, AND ECOSYSTEMS

II.1 TECHNICAL ENABLERS AND FUNCTIONAL SCOPE

Three things, one in hardware and two in software, came together between roughly 2023 and 2025 to make on-device generative AI something you could actually ship. On the hardware side, the neural processing units in phones now push through trillions of operations a second. The software half is two-fold: compression tricks, quantisation and distillation and pruning, shrink a model's parameter count without wrecking its accuracy, and the operating systems opened up APIs that let an assistant reach into other apps through structured intents, what Apple calls App Intents and Android calls App Actions.⁵

For a shopper, all of that technical detail just shows up as a list of things the assistant can suddenly do. It can read a photo of a product or a screenshot of a price tag. It can hold the thread of a conversation across several turns. It can reach the preferences you have saved, fire off actions inside installed apps like adding to a cart or tracking a parcel, and keep at least some of this working when the signal is bad. Each of those lines up with a piece of the buying process, which is what the next section sets out.




Figure II.1 — Assistant-mediated commerce ecosystem (conceptual architecture)

Figure II.1 sketches the stack a purchase actually travels through. At the top is the operating system, holding together the local neural hardware that runs the model, the conversational interface that hears the request, and the tokenised payment rail that handles the bank-grade authentication. Underneath that sit App Intents on Apple and App Actions on Android, the structured hooks (add to cart, track a shipment, redeem an offer) that a retailer's app has to implement if it wants to be part of any of this at all. Below those are the product and inventory feeds, the machine-readable lists of identifiers, attributes, stock and price; without them the assistant's answers fall apart no matter how clever the model is. Off to one side runs a hybrid cloud branch that takes over the heavy summarising, or slips in a sponsored placement, whenever the local

model runs out of road, and it drags the latency and privacy questions from earlier right back in. The thing to take away from the diagram is where the power sits vertically. The OS owner controls the assistant and the payment layers even when the goods themselves come from independent shops, which leaves merchants competing on feed quality and integration while the platforms compete on trust, speed, and who gets the default slot.

Who even gets to use this is gated by the hardware. Apple Intelligence, at launch, only ran on the newer iPhones with an A17 Pro chip or better, which left an enormous base of older phones out entirely (Apple Inc., 2024). That opens up a kind of digital divide inside a single ecosystem: the people who can afford to upgrade get assistant-led shopping first, and everyone else stays on the old browser route for a while longer. The practical upshot for a firm is that targeting the assistant means segmenting your audience by which generation of phone they are carrying, not just by the usual demographics.

II.2 A TAXONOMY OF PURCHASING-RELATED ASSISTANT FUNCTIONS

The assistant pokes into every stage of a purchase, just not to the same depth at each. Right at need recognition it can throw up a contextual reminder, the classic "remind me to buy nappies when I'm near a supermarket", or a trend prompt designed to stir demand you did not know you had. Once you actually start searching, it handles the constraint-style queries ("wireless earbuds under €80 with multipoint pairing") and the visual lookups that turn a snapshot of a shoe into a shortlist, both of which cut the cost of looking. At the evaluation stage it summarises reviews when asked ("what do people complain about with the battery?") and lines options up side by side, doing the compressing you would otherwise do in your own head. When it comes to buying, it

deep-links you into checkout ("buy the last one in my size from Zara") and settles up through a tokenised wallet. And afterwards it tracks the order and gets returns going. The depth is the part that really varies. Search and evaluation are where the generative model earns its keep. Buying mostly just hands off to app infrastructure that already exists. And the after-sales side depends far more on how widely the assistant has been wired in than on how sophisticated the underlying model is.

II.3 ADOPTION DETERMINANTS AND USER TRUST

If you take UTAUT and bolt privacy onto it, three things look like they really decide whether someone shops through an assistant, as opposed to just asking it for the weather. The first is plain usefulness. McLean and Osei-Frimpong (2019) found that the strongest reason people reach for a voice assistant is simply that it is useful, with privacy worry tugging the other way, and now that the summaries are getting genuinely accurate, that usefulness may start to reach into mid-involvement categories it never touched before. Privacy reassurance is the second, and it bites hardest in sensitive moments, a health product, or a gift you would rather nobody saw coming; Awad and Krishnan (2006) showed just how much privacy concern governs how far people will accept personalisation, and keeping the work local can tilt that balance, but only if people actually grasp what still leaves the phone. The third is ease. Talking is a lower bar than learning an app's menus, which helps occasional and older users in particular, though my hunch is that knowing how to phrase a good prompt just becomes the new skill that sorts people out. Underneath all three sits trust, and trust is brittle. A single hallucinated spec, or one invented review summary, can put someone off delegating for a long time (Ji et al., 2023), and because pushing quantisation too hard on the device

can raise the error rate, firms are stuck with a real accuracy-versus-speed trade-off that runs straight through to returns and brand damage.

II.4 COMPARATIVE ECOSYSTEM ANALYSIS: APPLE, GOOGLE, AND SAMSUNG

Apple leans hardest on privacy as the selling point. Its assistant does as much as it can on the phone itself and pushes only the overflow out to what Apple calls Private Cloud Compute (Apple Security Research, 2024), with shopping running through App Intents into whatever retailer apps you happen to have installed and out the other side through Apple Pay. The walled garden is rather the whole point here. Because the assistant's memory of you and your saved cards only work inside Apple's world, leaving gets more expensive the longer you stay. I notice this on my own iPhone: when I ask Siri to find something, it tends to steer me into Apple's own apps and the App Store before anything else, and once Apple Pay is set up it is easily the least-effort way to check out, which is the self-preferencing tendency in miniature. The thing Apple advertises is that the assistant can read your messages, calendar and files locally to figure out what you need, without that raw content ever reaching anyone else. For a retailer, getting into this channel means adopting App Intents and clearing Apple's privacy-label rules, and those are fixed costs that a big chain shrugs off far more easily than a small shop can.

Google's Android is the looser and messier ecosystem. It allows much broader app integration, but the experience fragments across every manufacturer's own version of Android, and Gemini leans on Google's existing search and shopping graph, which is exactly where the self-preferencing worry under the DMA begins, given that advertising tied to search still supplies the overwhelming majority of Alphabet's revenue (Alphabet Inc., 2024). The technical split is that Gemini Nano does the lighter classification and

extraction work on the device and sends the harder queries off to the cloud (Google DeepMind, 2024). Since Google is already tagged as a gatekeeper for both Search and Android, an assistant answer that slips Google Shopping above a merchant's own link walks straight into the same kind of scrutiny as the old EU Google Shopping case. Samsung is the interesting in-between case. It runs largely on Google's underlying models but layers its own features on top, things like live translation, and how much of the commerce side is actually switched on depends heavily on the region and the carrier. Its sheer hardware volume makes it the case that matters most outside the US and the EU, even though the public documentation is noticeably thinner. Inside the EU there is an odd duplication: Samsung Pay and Galaxy AI sit right next to Google Pay and Gemini on the very same phone, and I suspect plenty of users genuinely lose track of which assistant is even handling a given request. That confusion carries a cost of its own, because some people will just give up and stop delegating at all.

Table II.1 lines the three ecosystems up against each other on the dimensions that matter for shopping.

Table II.1 — Simplified ecosystem comparison

Dimension	Apple	Google/Android	Samsung
Primary on-device AI branding	Apple Intelligence	Gemini Nano	Galaxy AI
Payment rail	Apple Pay (closed)	Google Pay (semi-open)	Samsung Pay (regional)
App integration model	App Intents (curated)	App Actions (broader)	Mixed
Privacy narrative	Strong local emphasis	Hybrid; search integration	Follows Google + hardware layer
Lock-in intensity	High	Medium–high	Medium

None of the three publishes anything like an "assistant-attributed" sales figure, so there is simply no clean number to point to. The outside estimates that do exist put voice and conversational commerce at low single-digit percentages of total e-commerce (OC&C Strategy Consultants, 2018), and while that share has clearly grown since, it is still

small. So instead of waving around a market-share figure I could not really stand behind, I stay with how the mechanisms work and where they seem to be heading.

II.5 INTEGRATION DEPTH AND MARKET PRACTICE

What actually decides the commercial impact is how deeply the assistant is integrated, not how big the underlying model is, and you can watch the same pattern play out across a few different settings. Take replenishment. Amazon let you reorder through Alexa by voice years ago, and smartphone assistants now copy that through retailer App Intents; "reorder my usual detergent" folds a whole multi-step decision down into a single habit, and it shrinks the set of options you consider to exactly one, which rewards whichever shop already holds your purchase history. In a physical store it looks different again. Google Lens and Apple's Visual Intelligence let you compare products in the aisle just by pointing the camera, and because the image parsing can run locally it keeps working even on a weak in-store connection, which turns the screws on showrooming as the cost of comparing in the aisle falls to almost nothing. Then there are the mid-involvement durables. The electronics chains that built their own spec-comparison tools watch that advantage evaporate the moment the OS assistant does the same job for free, and they tend to fight back with exclusive stock, bundled services, or loyalty perks you can only see inside their own app. If there is a thread running through all of this, it is that assistants squeeze the path for low- and mid-involvement goods first and leave the high-involvement ones partly shielded, at least until accuracy and the liability rules catch up.

II.6 THE EUROPEAN AND ITALIAN RETAIL CONTEXT

Europe has some of the highest smartphone penetration anywhere, and Eurostat (2024) shows e-commerce participation still climbing, with mobile the default channel for the younger crowd. Italy is a useful case precisely because its retail is such a mix: the big national grocery chains, the international fashion marketplaces, the electronics specialists, and the assistant meets each of them differently. Grocery is the most obvious fit. Esselunga's app already pulls loyalty, shopping lists and targeted promotions into one place (Esselunga SpA, 2024),⁸ and an assistant that could read your reorder patterns locally, without shipping your basket history off to cloud advertisers, might add a geofenced nudge ("collect your order near Ancona within the next 30 minutes") that tips planned buying over aimless browsing, though only where people genuinely trust the local data handling. Fashion is a different story. Zalando runs at huge scale in Italy on free returns and sizing tools, and those returns are a well-documented cost centre (Zalando SE, 2023),⁷ so visual search cuts the cost of looking while generative sizing advice can do the exact opposite, raising the odds of a bad fit and feeding the returns pile. In electronics the in-aisle visual search makes showrooming worse, and a chain like MediaWorld pushes back with price guarantees, bundled installation, or stock you simply cannot get elsewhere (MediaWorld, 2024).

The regulation here sits at two levels at once, the national and the European. Italy's competition authority keeps an eye on platform conduct alongside the DMA enforcement (Autorità Garante della Concorrenza e del Mercato, 2020), and an assistant that leans toward platform-affiliated merchants over Italian SMEs could end up answering to both. What the Italian case really brings out is that the shape of local retail still steers how assistant commerce unfolds even when the operating systems themselves are global, and that the willingness to delegate tracks age and prior

smartphone experience (ISTAT, 2024), which means the gains are going to land unevenly across different groups.

CHAPTER III — IMPACT ON CONSUMER PURCHASING DECISION- MAKING

III.1 NEED RECOGNITION AND PROBLEM AWARENESS

Problem recognition is the moment something flips you from not-shopping to shopping, whether the trigger comes from inside you or from the world around you (Solomon, 2020). Phones were already good at supplying the outside trigger, through notifications and targeted ads. What an assistant adds is a trigger built from what it knows about your own life: a calendar that says your partner's birthday is tomorrow, a location that says you are walking past a shop holding something on your wish list, a habit that says you tend to run out of coffee pods about every six weeks.

There is a real upside to this. A nudge that turns a future need into a calm purchase now can save you from the panicked last-minute buy at a bad price, at least if you are the sort of person whose preferences stay stable over time. The catch is sponsorship. Once those reminders are paid for, they can start manufacturing purchases that would never have happened on their own, and at that point they look a lot more like persuasive advertising than like helpful information (Armstrong, 2008).

Keeping the processing on the device matters most right here, because the signals involved, health data, the contents of your messages, your photos, are about as sensitive as personal data gets. Someone who has firmly switched cloud tracking off might still let the phone work things out locally and prompt them. And whoever lands the deal with the phone maker for the "suggested reorder" slot gets first access to you, before a

rival has even made it into the running. It is the digital version of prime shelf placement in a supermarket, only far more personalised and much harder to spot.

In Italy the grocery apps already bombard you with promotions, and an assistant sitting on top could go one of two ways. It could pull those reminders together across the shops whose apps you have linked, which would be genuinely useful. Or it could funnel all the demand to whichever chain wins the default integration first, which would shut out the smaller chains that cannot spare a developer. My own expectation is more planned and replenishment buying overall, with a bit more impulse buying wherever the prompts are well-timed. The brands that expose clean subscription or reorder APIs are the ones set up to gain.

III.2 INFORMATION SEARCH AND CONSIDERATION-SET FORMATION

Search is where the assistant cuts the old Stiglerian cost most visibly. One conversational question can gather up the attributes, the ratings and the prices that used to mean clicking through retailer sites and review platforms one after another.

What it does to the size of your consideration set is genuinely two-sided. On one hand it compresses things: you get three to five options and you satisfice from those (Simon, 1955). On the other it can broaden them, because visual and multimodal search can throw up things you would never have thought to type, a niche label or a seller in another country. Which way it tips seems to ride on how specific your question is. Ask for "good headphones" and you get a short, popularity-skewed list, which is exactly what the recommender-system research would predict (Ricci et al., 2015). Ask for "ANC headphones under 200g, USB-C, with replaceable ear pads" and the assistant can actually widen your options, matching combinations a generic bestseller list never bothers with.

Speed feeds back into how hard people search. When asking one more question costs almost nothing, you tend to keep asking rather than settle, which claws back some of what the compression took. Hybrid setups add their own snag here: if the assistant has to fall back to the cloud it needs a connection, whereas the on-device functions keep working even on bad signal, like reading a price tag through the camera in a shop where a cloud-only assistant would just stall.

There is a clear asymmetry of trust running through all of this. People hand over the low-stakes searches long before the high-stakes ones, which is just what McLean and Osei-Frimpong (2019) found. And summarising reviews carries its own risk, because the model can drop a recurring complaint or invent a consensus that was never there (Ji et al., 2023). Strictly, a careful shopper ought to distrust any summary that does not link back to the actual reviews, but I have seen no sign that mainstream users behave that way yet.

It helps to trace one example all the way through. Say you are in Zalando's app looking at running shoes, and you lean on the assistant's review summary before settling on a size. If the local model squeezes those reviews down without flagging the fit complaints that drive returns, you have handed the evaluation to the agent, and you may well end up feeding the return pile Zalando already manages on an enormous scale (Zalando SE, 2023).⁷ An omission back at the search stage turns into a cost at the post-purchase stage, which is precisely the hallucination-and-omission problem the summarisation literature keeps warning about (Ji et al., 2023). The search itself definitely gets shorter. Whether your set of options gets wider or narrower, honestly, is not clear to me. What does look certain is that traffic drifts off the standalone comparison sites and toward the assistant,

and that clean, well-structured product data becomes a real asset for whichever merchant happens to own it.

III.3 EVALUATION OF ALTERNATIVES AND HEURISTIC SUBSTITUTION

Evaluation is the weighing-up: price against quality against ethics against looks, almost always under time pressure. Because attention is scarce, people fall back on shortcuts, the brand they already trust, the star rating, the reviewer whose taste they share (Tversky and Kahneman, 1974). An assistant automates those shortcuts for you, handing over a pros-and-cons list, a declared "winner", a tidy table of specs.

What is really going on is that the shortcut gets outsourced to an algorithm. That saves you the mental effort, but it also hands the agent the authority to judge. Thaler and Sunstein (2008) made the point that the way choices are arranged, the defaults, the framing, the order they arrive in, steers people without ever removing an option. An assistant that lists Option A above Option B, or describes A in warmer language, is doing exactly that, even while it presents itself as neutral.

On-device personalisation pours in another ingredient. It learns from your own history, the things you bought, the suggestions you waved away, the trails you left, all of it kept on the phone. That can feel less invasive than the old cookie-based tracking across the whole web, since there is no obvious cross-site surveillance, and that feeling genuinely matters for whether people accept it. But the preference model still bakes in path dependence. Your early choices tilt what you get shown later, which can harden brand loyalty well past the point a coldly rational shopper would keep updating (Shapiro and Varian, 1999).

For the big, high-involvement purchases, electronics, anything regulated like a financial product, people still hold back from delegating the evaluation. They want the primary

sources, the warranty, the disclosures a breezy summary tends to chop off. So decisions get faster across the board, but they also get more variable in how well they match what the buyer actually wanted. Defaults and whatever sits at the top of the list take a bigger share, and the established brands gain ground, unless the assistant has been deliberately tuned to reward variety or a lower price.

III.4 PURCHASE EXECUTION AND PAYMENT INTEGRATION

Buying, in the narrow sense, is choosing how to pay, authorising it, and getting the confirmation. Assistants shorten the run from wanting something to actually paying for it, through deep links and a built-in wallet. Apple Pay and Google Pay keep your card details as tokens, and the fingerprint or face check on the device is enough to satisfy the strong-customer-authentication rules PSD2 imposes in the EU (European Banking Authority, 2019).⁶

There is an old idea in transaction-cost economics, going back to Coase (1937) and dragged into the digital world by Bakos (1997), that when the cost of coordinating a deal falls, more deals get done. One-tap, assistant-started checkout strips out exactly the friction that used to kill carts, especially on the mobile web where people give up halfway through filling in a form (Häubl and Trifts, 2000). The conversion bump from removing checkout friction is one of the better-documented effects in the whole e-commerce literature; the assistant just carries that logic over to starting the purchase by talking.

In Italy wallet use still varies a lot by age and by which shops will take it, but the supermarket chains and the marketplaces are accepting Apple Pay and Google Pay more and more. An assistant-started checkout that ends in a tokenised wallet payment lines up

neatly with the EU authentication rules while sparing you from re-typing card details, which is a tidy example of cutting friction and staying compliant at the same time. The risk all bunches up at the authorisation step. A misheard or misread command, "buy" against "why", can fire off an order you never meant, and accidental purchases by children or a phone left in a pocket were a real headache in the early days of voice commerce. An on-device confirmation screen helps, but it puts a little of the friction back, and that is the genuine tension here: every bit of seamlessness you add chips away at the protection.

The platform's payment rail also rakes in the interchange fees and the transaction data, which deepens the lock-in, and a merchant who takes the platform wallet is trading lower friction for dependency, something Chapter IV comes back to. So I would expect conversion to rise, most visibly on cheap items, and wallet payment to take a growing share of transactions. The price for all that convenience shows up as more disputes, the times a misread command produces an order the buyer never actually wanted.

III.5 POST-PURCHASE BEHAVIOUR, LOYALTY, AND RETURNS

After the purchase comes everything else: whether you were happy, whether you buy again, the review you leave, the thing you send back, what you tell your friends. The assistant pitches in with tracking the order, nudging you about delivery, and walking you through a return. Done well, that trims the merchant's customer-service costs and the buyer's anxiety at the same time, which is a real win for both sides.

Returns are where the money really is. If the assistant fed you bad information back at the search or evaluation stage, the wrong size, a compatibility claim that turned out to be false, you are more likely to send the thing back, and in categories with free returns that eats straight into the margin. Zalando and Amazon have swallowed return logistics

as a cost of doing business; a small merchant has not, and takes asymmetric damage from an assistant error it never had any control over.

Loyalty schemes plug into the assistant the moment the local memory holds your membership and applies the discount for you automatically. That convenience raises the cost of leaving, because walking away from the ecosystem now means losing all the context the assistant has built up about you, not just your accumulated points.

Word-of-mouth may drift from public social posts to private questions put to the assistant ("should I recommend this to my sister?"), which makes the sentiment far harder for marketers to see and knocks straight on into how well advertising works. The net picture is genuinely mixed. Service-contact costs clearly fall. The effect on returns could honestly go either way, since convenience and misinformation pull against each other. And switching costs climb as loyalty and memory get folded into the same place.

III.6 PULLING THE STAGE-LEVEL EFFECTS TOGETHER

Table III.1 — How on-device assistants affect each purchasing stage

Decision stage	Key assistant capability	Economic mechanism	Likely consumer outcome	Likely firm/market outcome
Need recognition	Contextual prompts, reorder suggestions	Attention reallocation; inter-temporal nudging	More timely replenishment; risk of unplanned purchases	OEM/partner control of demand timing
Information search	NL queries, visual search	Search-cost reduction	Faster discovery; possible omission of alternatives	Traffic diversion from independent comparators
Evaluation	Review synthesis, comparisons	Heuristic substitution; choice architecture	Faster decisions; default bias	Incumbent advantage; sponsored placement rents
Purchase	Deep links, wallet integration	Transaction-cost reduction	Higher completion rates	Platform capture of payment data and fees
Post-purchase	Tracking, returns, loyalty hooks	Service-cost reduction; lock-in	Convenience; path-dependent loyalty	Margin pressure if returns rise

Reading across Table III.1, the assistant's reach is wide but distinctly uneven. It is strongest where you search, evaluate and pay, and much more conditional at the two ends, need recognition and after-sales, where everything hinges on how deeply it has been wired in. The same four mechanisms keep resurfacing: search costs fall, the agent

takes over the heuristics, its framing turns into a kind of choice architecture, and switching costs compound over time.

III.7 WELFARE AND INTER-STAGE DYNAMICS

Convenience on its own is a poor way to judge whether any of this is actually good for people; you have to set the gains and the losses next to each other. Table III.2 does that, one channel at a time.

Table III.2 — Welfare gains and losses, by channel

Welfare channel	Potential gain	Potential loss
Search	Time savings; better matches	Omitted alternatives; sponsored bias
Evaluation	Cognitive relief	Systematic mis-ranking; hallucination
Price	Lower dispersion if comparison is honest	Hidden fees; personalised mark-ups
Privacy	Local processing reduces exposure	Hybrid cloud leakage; opaque consent
Competition	Faster entry for merchants with open APIs	Foreclosure of comparators; self-preferencing
Autonomy	Easier opt-out than high-pressure sales	Default bias; less deliberate comparison

Which side comes out ahead depends almost entirely on the category and on how the thing is built. A replenishment order that is accurate and honestly priced almost certainly leaves the shopper better off. A high-involvement purchase decided on the back of an unchecked summary can leave a careful buyer worse off than if they had done the legwork themselves, even while someone short on time and happy to satisfice still gains. That is why I do not think it makes much sense to call assistant commerce simply good or bad for consumers. The useful questions are narrower. Can you tell when something is sponsored? Can you take your data and walk? Can anyone outside the company audit how the rankings get made? Get those three right and the gains tend to win out.

The stages also loop back on each other in a way a one-shot model misses. A return caused by a bad evaluation pushes the shopper into heavier searching next time, and it chips away at their willingness to delegate at all. A smooth, frictionless purchase does the reverse: it hardens the habit and shrinks the next consideration set through loyalty memory. So your early experiences end up setting the long-run pattern and deepening the lock-in. There is a second loop running from need recognition back to data. Every contextual prompt that nudges a trial purchase sharpens the on-device profile, which makes the next prompt land better, and leaves any newcomer stuck with the cold-start problem recommender systems have always struggled with (Ricci et al., 2015).

Regulation can break into these loops. A sponsorship label at the evaluation stage might cost a few conversions today while keeping the comparison sites alive for tomorrow, which is exactly the short-term-pain-for-longer-term-competition trade-off that Armstrong (2008) argues competition and consumer authorities have to weigh together rather than separately.

CHAPTER IV — ECONOMIC AND STRATEGIC IMPLICATIONS

IV.1 MARKET STRUCTURE AND ECOSYSTEM LOCK-IN

Once an assistant lives inside the operating system, the company that owns that system ends up holding three things at the same time: the assistant you actually talk to, the wallet you pay with, and the gate every app has to pass through to get installed. That bundle is the real story of this chapter. If the assistant quietly steers you toward the platform's own marketplace or its preferred payment option, that is self-preferencing in everything but name, and it is the behaviour the Digital Markets Act was written to catch (European Commission, 2022).

The firms that lose out are easy to overlook, because nothing dramatic happens to them. Independent price-comparison sites and small marketplaces do not get blocked; they just stop being visited. When the assistant answers the question itself instead of handing you a page of links, the click that used to land on a comparison site never happens. The web already gave us a preview of this: as Google began answering more queries on the results page itself, a majority of searches started ending without any onward click, pulling traffic away from the very publishers being summarised (Fishkin, 2019). An assistant pushes the logic further, since there is not even a results page left to scroll through.

Why is it so hard to walk away from one of these ecosystems? Some of it is the obvious stuff, saved cards and downloaded apps. What I find more important is the assistant's memory of you. It learns what you buy, what you turn down, even how you phrase a request, and none of that travels with you. You cannot export your "history with the assistant" to a rival phone the way you copy contacts. Shapiro and Varian (1999) made the classic argument that lock-in rents fade once switching costs drop, but here I think the opposite happens. The longer you use the assistant, the better it knows you, so the cost of leaving grows rather than shrinks. It is learning-by-using, except the platform keeps most of the payoff.

Add all of this up and the advantage drifts toward whoever owns the platform, plus the handful of large retailers that can afford to integrate deeply. Smaller sellers are not really punished. They just risk going invisible inside a shortlist the customer never assembles by hand.

IV.2 IMPLICATIONS FOR RETAILERS AND CONSUMER BRANDS

For years the retail game was about being found by Google. Then it shifted to being placed well on a marketplace like Amazon. The next version of that same game is about being legible to an assistant, which for lack of a better term I will call assistant optimisation.⁹ The old SEO playbook ran on keywords, backlinks, and chasing a high rank, and you judged it by how many people clicked through. This new version is duller and far more technical. It rewards clean, machine-readable product data: sensible attributes, stock counts that are actually current, a standard product code like a GTIN, and intent handlers so that when someone says "add this to my cart" the assistant knows exactly which item is meant. There is no penalty for getting it wrong. You simply never appear. And the worst case is structural, because the assistant can favour its own shop over yours.

This divides retailers sharply. Amazon, Walmart, or Zalando can put a team on assistant and voice integration tomorrow. A small merchant usually cannot. The businesses I worry about most are the kind that are everywhere in central Italy, where I am writing this: regional food producers, the leather and craft workshops you find around the Marche, firms whose products are genuinely good but who have nobody on staff to wire up App Actions. If the assistant only surfaces merchants that have done the technical work, those producers can fall out of the shortlist completely. Not because a customer rejected them, but because no customer ever saw them.

Brands face a related decision. If an assistant can deep-link cleanly into a brand's own app, that brand can push direct sales and skip the marketplace middleman. A brand that lives on someone else's marketplace gets squeezed twice: once by the platform's fees, and now by the uncertainty of whether the assistant will mention it at all. There is also a quieter problem that brands are only starting to notice. When an assistant summarises

reviews and tells the shopper "what people think of this brand," it is rewriting the brand's story in its own words, and the careful positioning written for a human reader rarely survives that compression. Learning to manage how a machine paraphrases you is turning into part of the marketing job.

IV.3 PRICING, TRANSPARENCY, AND COMPETITIVE DYNAMICS

In theory, cheaper search should sharpen price competition, because once you can see who is overcharging, the overcharging gets harder to sustain (Stigler, 1961). An assistant that genuinely quotes the lowest price across the shops it is plugged into ought to squeeze margins. But it might not play out that way. If the assistant leads with total cost, folding in delivery and the returns policy and how reliable the seller is, or if it just keeps steering you back to the merchants it already trusts, prices can stay stubbornly spread out in a way that is not really competitive at all.

The other side of pricing is discrimination. Dynamic pricing and personalised offers, fed by the loyalty data sitting on the device, make much finer price discrimination possible. Charging different people different prices off protected characteristics is still illegal in plenty of places, but sorting customers by behaviour, the frequent buyer against the occasional one, is completely routine. The twist is that an assistant with local memory can do that sorting without the obvious cookies regulators learned to watch for, which makes the whole thing harder to police.

And there is a plain transparency gap: most of the time you cannot tell whether a recommendation is organic or paid for. The FTC's old guidance on disclosing digital advertising does technically cover a conversational interface (Federal Trade Commission, 2013), but enforcing it on a tiny screen, or in a voice-only exchange with no screen at all, lags well behind what we have come to expect on the web.

IV.4 MACROECONOMIC CHANNELS AND AGGREGATE CONSUMPTION

The macro picture deserves real caution, so everything here is qualitative; I am sketching channels, not putting numbers on them. Easier replenishment might just smooth out how people buy staples without changing what they spend over a year, while a well-placed impulse prompt can pull discretionary spending forward in time. The hours saved on dull, routine buying could get redirected toward services and experiences, or they might simply mean less time spent shopping at all, a small household productivity gain not unlike what self-checkout once did for supermarkets. And if assistants genuinely improve how well buyers are matched to products, the advertising spend needed per sale can fall, which reshuffles marketing budgets without doing much to GDP.

A couple of the channels are subtler. If sales keep concentrating on the top products and the biggest platforms (Ricci et al., 2015), that could thin out the entry of small niche players, with effects on innovation that are genuinely hard to pin down. The one I find most interesting is measurement. National accounts log what was bought and through which channel, but not whether the purchase began life as a conversation, so a growing slice of assistant-driven sales may just get filed under generic "e-commerce growth" until the disclosure standards catch up, and a central bank could miss the small, prompt-driven demand bumps buried in the aggregate series. Whether the freed-up time turns into more buying, more leisure, or more work is genuinely open, and it carries a distributional edge, because it is the time-poor, higher-income households that tend to value the saving most.

IV.5 REGULATORY AND CONSUMER-PROTECTION CONSIDERATIONS

Regulators are going to run into assistant-mediated commerce on three fronts, and it is easier to walk through each with a concrete case than in the abstract.

The first is competition. The DMA's obligations on gatekeepers already bar tying and self-preferencing, so picture this: you ask the OS assistant to "find the cheapest washing machine with delivery to Ancona," and the top result it hands back is a retailer that pays into the platform's own checkout and advertising stack, while a cheaper local appliance dealer who never implemented the platform's payment SDK sits three results down, or gets summarised away entirely. Nothing in the answer is false, yet the ranking still favours the gatekeeper's commercial partners. That is exactly the conduct Article 6 of the DMA is aimed at, and it would be very hard to catch without the kind of audit access regulators do not currently have. Interoperability rules may eventually force assistants to show third-party merchants on comparable terms, which is as much an enforcement headache as a technical one.

The second is consumer protection, where the hallucinations turn into liability questions. Suppose the assistant tells you a phone charger is "compatible with your laptop" when it is not, and you buy it on that basis; or it calls a supplement "gluten-free" after misreading the product page. Who is on the hook, the company that ships the assistant, the firm whose model produced the claim, or the merchant whose listing got misread? The law has no clean answer yet, and the revised EU Product Liability Directive, which expressly brings software and AI systems within the scope of no-fault product liability (European Commission, 2024), is a first step toward sorting out who carries that responsibility.

The third is privacy. Doing the work on the device sits comfortably with the GDPR's data-minimisation principle (European Union, 2016), but the hybrid fallback muddies it.

If a question about an intimate health product is parsed locally and then escalated to the cloud for a fuller answer, the redacted-but-still-sensitive bit that leaves the phone needs a documented legal basis to travel, and most people will have no idea the hand-off even happened. Disclosing sponsored content hits the same perception problem: it only really satisfies unfair-commercial-practices law if the user can actually tell a recommendation was paid for, which is far harder to convey in a single spoken line than on a web page. The way Italy and the wider EU bring in the AI Act (Regulation (EU) 2024/1689) adds another layer, sorting systems by risk so that general-purpose models and high-risk uses pick up real documentation and oversight duties.¹⁰ A shopping assistant that strays into advising on product safety or anything health-adjacent may eventually face stiffer conformity checks than one that just recommends films, and that compliance cost is something the OEMs will push down through their developer programmes. There is also every chance the AGCM and the European Commission end up coordinating on cases where assistant defaults nudge people toward gatekeeper marketplaces, especially where Italian SMEs depend on being visible in aggregator apps (Autorità Garante della Concorrenza e del Mercato, 2020). National competence does not vanish under the DMA; it sits alongside the EU gatekeeper rules for conduct that harms competition at home.

IV.6 WHAT EACH ACTOR SHOULD DO

Table IV.1 distils the discussion into the priority moves for each actor and the cost of standing still.

Table IV.1 — Priorities and risks by actor

Actor	Priority actions	Risk if inactive
OEM / platform	Publish transparency on sponsored recommendations; open APIs for third-party merchants; accuracy auditing	Antitrust fines; reputational damage from hallucinations

Actor	Priority actions	Risk if inactive
Retailer / brand	Structured product data; App Intents/Actions; monitor generative brand representation	Visibility loss; margin compression
Independent comparator	Pivot to verification, open data, regulatory advocacy	Traffic collapse
Policymaker	Enforce disclosure; clarify liability; monitor self-preferencing	Consumer harm; market foreclosure
Consumer advocacy	Literacy programmes on delegation limits	Over-trust in summaries

The split of responsibilities is reasonably clear. Firms do the technical integration, and policymakers supply the transparency and interoperability mandates that keep it honest. Consumers come out ahead when search and transaction costs fall and nobody is hiding the sponsorship or serving up errors.

IV.7 REGIONAL POLICY, SMES, AND THE ITALIAN MARCHE CONTEXT

The distributional side of all this matters a great deal for regional economies built on small firms. Italy is a good example, since SMEs employ most of the non-financial workforce, and a great many of them have no in-house team to build App Intents or keep a structured feed running. When the assistant favours merchants with full digital integration, the gap between places can widen: the well-served firms in Milan or Rome gain visibility, while an artisan producer in the Marche or Abruzzo stays reachable only through an aggregator that takes a commission on every sale. There are tools that could help, the EU cohesion funds, the PNRR digitalisation vouchers, training run through the chambers of commerce, and they could pay for "assistant-readiness" audits that check whether a firm's products even carry the right identifiers, whether its intent handlers actually work, and how its brand comes across in a generative summary. It is not so different from the export-readiness programmes that help firms win shelf space abroad. The Marche is a particular case, because so much of its retail is tied to tourism, wine, leather, food, and it mixes high experiential involvement with one-off purchases by

visitors passing through. So I would expect the assistant's real impact to land less on converting a first-time boutique sale and more on the repeat reorder from a visitor who got home and wanted more, which suggests support should follow category-specific paths rather than one blanket "AI commerce" template. Further out, the DMA's interoperability obligations might force the gatekeepers to allow third-party assistants or payment services, and that could finally let Italian merchants route customers through a domestic banking-app wallet instead of a US platform's rails, loosening the lock-in at least a little.

CHAPTER V — DISCUSSION, LIMITATIONS, AND RESEARCH AGENDA

Synthesis of findings

If there is one thing to carry away from all this, it is that an on-device assistant is closer to an economic intermediary than to a neat extra feature on a phone. It has a hand in every step of how people shop, from spotting a need to chasing up a late delivery, but its grip is far from even. Most of the real work happens through cheaper search and through the way it takes over the comparing, while the lock-in builds more slowly, through payment and through everything the assistant remembers about you. Set against the old cloud-only assistants, the local design rewrites the privacy-versus-personalisation bargain and can pull in people who used to keep their distance. And next to ordinary app shopping, it shortens the whole journey and drags market power toward whoever owns the operating system.

What struck me most in comparing the three is how little they really differ where it counts. Apple, Google and Samsung have all landed on the same playbook: a wallet built in, developer APIs for triggering actions, generative summaries of products. What separates them is mostly the story they tell about privacy, how exposed they are as

gatekeepers, and how tightly they run their app stores, not any difference in appetite for making money. If I had to hand a practitioner a horizon, I would say two to three years. Assistant-attributed sales will stay hard to measure for a while yet, but how deeply you integrate with the APIs is already starting to decide who shows up and who does not.

Limitations

All of this comes with real limits, and I would rather be upfront about them than bury them. There is still very little peer-reviewed work that actually measures on-device assistant shopping, so a good deal of my reasoning leans on neighbouring research into voice commerce, chatbots and recommenders. The companies' own documentation oversells the benefits, and reading it sceptically only takes you so far. The technology and the rules are both moving fast enough that some of the specifics here will be out of date before long. Sticking to EU and US ecosystems also means the picture may not hold up in places where mobile money is the norm or privacy expectations are simply different. And the big one: the only primary evidence I gathered is a small, self-selected survey ($n = 40$, Appendix C), which is suggestive at best and nowhere near a controlled test at the level of a real purchase. The map in Chapter III is still a structured set of hypotheses rather than a proven model; the survey lines up with it, but a convenience sample that size can do little more than hint. It is there mainly so that someone can begin checking the framework against real people.

Future research agenda

So where should this go next? The obvious move is to actually run the experiments. Compare conversion and return rates when a purchase starts with the assistant against when it starts in a normal app. Audit how neutral the recommendations really are and whether sponsorship is ever disclosed. Track over time how people's willingness to

delegate drifts from one category to another. On the macro side, someone will eventually be able to follow assistant-attributed sales once the platforms start reporting them at all. A much smaller version of this is doable in a single term, which is what the survey in Appendix C does on a convenience sample: it asks smartphone users which steps they would actually hand over, how far they trust local processing against the cloud, and where they draw the line completely. One methodological plea runs through all of it. Keep on-device, hybrid and cloud setups apart, because lumping them together hides the very privacy and latency differences the whole report turns on.

CONCLUSIONS

This report set out to understand how on-device assistants change the way people buy things on their phones, and what that means for competition, for company strategy, and for regulation. I used the old Engel–Blackwell–Miniard five-stage picture of buying as a backbone and leaned on search-cost theory, bounded rationality, choice architecture, and platform economics to read each stage. What came out of that is the mechanism map in Chapter III, which ties what an assistant can do to the economic effect it produces.

Here is the short version. The assistant matters most in the middle of the journey, when you are searching, comparing, and paying, and matters less, or only sometimes, at the very start and the very end. It works through the same few levers every time. Searching gets cheaper. The job of comparing options gets handed to the agent. The way that agent lays out the choices steers them. And walking away gets harder as your payment details and your personal history pile up on the device. Power tends to flow to whoever owns the operating system and to the large retailers that can afford to plug into it, which leaves comparison sites and small shops exposed unless the rules on interoperability and

transparency get tighter. So the sensible move for a firm is to treat assistant integration as real distribution work and to keep an eye on how its brand gets described, while regulators have to force sponsorship disclosure, settle who is liable when the advice is wrong, and apply the gatekeeper rules when an assistant quietly favours its own services.

I want to be careful not to oversell the word "on-device". Keeping the computation local really does cut some privacy exposure, but it guarantees nothing on its own, and it does little to stop the platform gathering power around itself. Whether any of this leaves shoppers better off comes down to unglamorous details: how accurate the assistant is, how neutral, and whether you can pick up your data and leave. Some of that is an engineering question and some of it is political.

For a business the practical takeaway is blunt. The assistant layer is becoming part of the distribution channel, and treating it as optional means risking a quiet demotion inside a shortlist the customer never builds by hand. For policy, the point is that the tools we already have, competition law and consumer protection, mostly still apply. They just need adapting to an intermediary that talks back instead of listing links, and that runs partly on the device itself.

There is one last irony worth sitting with. The single biggest purchase in this whole story, the phone itself, is what decides which assistant you get, and that assistant then shapes years of smaller purchases afterwards. That loop is the lock-in argument of this report in miniature, and it is why I think the subject deserves more serious attention than a phone feature usually gets.

Managerial and policy recommendations (summary)

A few concrete things follow for managers. Treat assistant integration like any other distribution channel, which above all means measuring it: how often your intent actually gets invoked, what converts when it does (where you can even track that), how often assistant-referred orders come back, and, every so often, how your brand is being described inside the generative summaries. For policymakers, the move is to stretch the digital-fairness tools that already exist, the sponsored-content labels, the DMA's self-preferencing rules, the product-liability reform, onto conversational interfaces, rather than sit and wait for some bespoke "AI shopping" statute. And there is a job for consumer education too. People need a basic delegation literacy: a feel for when to go and check a primary source, how to actually read one of those hybrid privacy notices, and how to climb back out of a default recommendation loop when it stops serving them.

Closing remarks

I do not want to overstate any of this. But the technology really is new, and the economics, even if it is inferential rather than measured, is structured enough to take seriously. The stakes are real too, for ordinary shoppers, for Italian retailers, and for the regulators keeping an eye on the platforms. What I have not done is settle the questions that genuinely need longitudinal data. This report is meant as a baseline instead, a way of framing the problem that later empirical work, including the modest survey and audit experiments described above, can test and build on.

REFERENCES

Alphabet Inc. (2024) *Form 10-K for the fiscal year ended 31 December 2023*. Mountain View: Alphabet Inc.

- Appel, G., Grewal, L., Hadi, R. and Stephen, A.T. (2020) 'The future of social media in marketing', *Journal of the Academy of Marketing Science*, 48(1), pp. 79–95.
- Apple Inc. (2024) *Apple Intelligence*. Cupertino: Apple Inc. Available at: <https://www.apple.com/apple-intelligence/> (Accessed: 4 June 2026).
- Apple Security Research (2024) *Private Cloud Compute: a new frontier for AI privacy in the cloud*. Available at: <https://security.apple.com/blog/private-cloud-compute/> (Accessed: 4 June 2026).
- Armstrong, M. (2008) 'Interactions between competition and consumer policy', *Competition Policy International*, 4(1), pp. 97–107.
- Autorità Garante della Concorrenza e del Mercato (2020) *Indagine conoscitiva sui Big Data (IC53)*. Joint inquiry with AGCOM and the Garante per la Protezione dei Dati Personali. Rome: AGCM. Available at: <https://www.agcm.it/> (Accessed: 4 June 2026).
- Awad, N.F. and Krishnan, M.S. (2006) 'The personalization privacy paradox: an empirical evaluation of information transparency and the willingness to be profiled online for personalization', *MIS Quarterly*, 30(1), pp. 13–28.
- Bakos, J.Y. (1997) 'Reducing buyer search costs: implications for electronic marketplaces', *Management Science*, 43(12), pp. 1676–1692.
- Card, S.K., Moran, T.P. and Newell, A. (1983) *The psychology of human–computer interaction*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Coase, R.H. (1937) 'The nature of the firm', *Economica*, 4(16), pp. 386–405.
- Engel, J.F., Blackwell, R.D. and Miniard, P.W. (1995) *Consumer behavior*. 8th edn. Fort Worth: Dryden Press.

Esselunga SpA (2024) *Esselunga* [online]. Available at: <https://www.esselunga.it/>

(Accessed: 4 June 2026).

European Banking Authority (2019) *Regulatory technical standards on strong customer authentication and common and secure communication (EBA/RTS/2017/02)*.

Paris: EBA.

European Commission (2022) *Regulation (EU) 2022/1925 (Digital Markets Act)*.

Brussels: Official Journal of the European Union.

European Commission (2024) *Directive (EU) 2024/2853 on liability for defective products (revised Product Liability Directive)*. Brussels: Official Journal of the

European Union. Available at: <https://eur-lex.europa.eu/eli/dir/2024/2853/oj>

(Accessed: 4 June 2026).

European Union (2016) *Regulation (EU) 2016/679 (General Data Protection Regulation)*. Brussels: Official Journal of the European Union.

European Union (2024) *Regulation (EU) 2024/1689 (Artificial Intelligence Act)*.

Brussels: Official Journal of the European Union.

Eurostat (2024) *E-commerce statistics for individuals*. Luxembourg: Eurostat. Available at: <https://ec.europa.eu/eurostat/> (Accessed: 4 June 2026).

Federal Trade Commission (2013) *.com disclosures: how to make effective disclosures in digital advertising*. Washington, DC: FTC.

Fishkin, R. (2019) *Less than half of Google searches now result in a click*. SparkToro.

Available at: <https://sparktoro.com/blog/less-than-half-of-google-searches-now-result-in-a-click/> (Accessed: 4 June 2026).

- Google DeepMind (2024) *Gemini Nano: on-device generative AI*. Android Developers.
Available at: <https://developer.android.com/ai/gemini-nano> (Accessed: 4 June 2026).
- Grewal, D., Bart, Y., Spann, M. and Zubcsek, P.P. (2016) 'Mobile advertising: a framework and research agenda', *Journal of Interactive Marketing*, 34, pp. 3–14.
- Häubl, G. and Trifts, V. (2000) 'Consumer decision making in online shopping environments: the effects of interactive decision aids', *Marketing Science*, 19(1), pp. 4–21.
- ISTAT (2024) *Cittadini e ICT: uso di smartphone e commercio elettronico*. Rome: Istituto Nazionale di Statistica.
- Iyengar, S.S. and Lepper, M.R. (2000) 'When choice is demotivating: can one desire too much of a good thing?', *Journal of Personality and Social Psychology*, 79(6), pp. 995–1006.
- Ji, Z., Lee, N., Frieske, R., Yu, T., Su, D., Xu, Y., Ishii, E., Bang, Y.J., Madotto, A. and Fung, P. (2023) 'Survey of hallucination in natural language generation', *ACM Computing Surveys*, 55(12), pp. 1–38.
- Luo, X., Tong, S., Fang, Z. and Qu, Z. (2019) 'Frontiers: machines vs. humans: the impact of artificial intelligence chatbot disclosure on customer purchases', *Marketing Science*, 38(6), pp. 937–947.
- McLean, G. and Osei-Frimpong, K. (2019) 'Hey Alexa... examine the variables influencing the use of artificial intelligent in-home voice assistants', *Computers in Human Behavior*, 99, pp. 28–37.
- MediaWorld (2024) *MediaWorld* [online]. Available at: <https://www.mediaworld.it/> (Accessed: 4 June 2026).

- OC&C Strategy Consultants (2018) *The talking shop: the rise of voice commerce*. London: OC&C Strategy Consultants. Available at: <https://www.occstrategy.com/en/report/the-talking-shop/> (Accessed: 4 June 2026).
- Puntoni, S., Reczek, R.W., Giesler, M. and Botti, S. (2021) 'Consumers and artificial intelligence: an experiential perspective', *Journal of Marketing*, 85(1), pp. 131–151.
- Qualcomm Technologies (2024) *On-device AI*. Qualcomm AI Hub. Available at: <https://aihub.qualcomm.com/> (Accessed: 4 June 2026).
- Ricci, F., Rokach, L. and Shapira, B. (eds.) (2015) *Recommender systems handbook*. 2nd edn. New York: Springer.
- Rochet, J.-C. and Tirole, J. (2003) 'Platform competition in two-sided markets', *Journal of the European Economic Association*, 1(4), pp. 990–1029.
- Samsung Electronics (2024) *Human-centric, hybrid AI opens up new possibilities*. Samsung Global Newsroom. Available at: <https://news.samsung.com/global/human-centric-hybrid-ai-opens-up-new-possibilities> (Accessed: 4 June 2026).
- Shapiro, C. and Varian, H.R. (1999) *Information rules: a strategic guide to the network economy*. Boston: Harvard Business School Press.
- Simon, H.A. (1955) 'A behavioral model of rational choice', *Quarterly Journal of Economics*, 69(1), pp. 99–118.
- Solomon, M.R. (2020) *Consumer behavior: buying, having, and being*. 13th edn. Hoboken, NJ: Pearson.

- Stigler, G.J. (1961) 'The economics of information', *Journal of Political Economy*, 69(3), pp. 213–225.
- Thaler, R.H. and Sunstein, C.R. (2008) *Nudge: improving decisions about health, wealth, and happiness*. New Haven: Yale University Press.
- Tversky, A. and Kahneman, D. (1974) 'Judgment under uncertainty: heuristics and biases', *Science*, 185(4157), pp. 1124–1131.
- Venkatesh, V., Morris, M.G., Davis, G.B. and Davis, F.D. (2003) 'User acceptance of information technology: toward a unified view', *MIS Quarterly*, 27(3), pp. 425–478.
- Zalando SE (2023) *Annual report 2023*. Berlin: Zalando SE.

BIBLIOGRAPHY

Sources consulted for background but not directly cited in the text.

- Brynjolfsson, E. and McAfee, A. (2014) *The second machine age*. New York: W.W. Norton.
- Davenport, T., Guha, A., Grewal, D. and Bressgott, T. (2020) 'How artificial intelligence will change the future of marketing', *Journal of the Academy of Marketing Science*, 48(1), pp. 24–42.
- Kahneman, D. (2011) *Thinking, fast and slow*. New York: Farrar, Straus and Giroux.

FOOTNOTES

1. The research problem couples micro purchasing stages with meso market structure and macro consumption, following conventions in the industrial organisation of digital markets (Rochet and Tirole, 2003; Bakos, 1997).
2. Under Article 6 of Regulation (EU) 2022/1925, gatekeepers may not rank their own downstream services more favourably than rivals'; an assistant that surfaces the platform's own marketplace or payment service raises the same self-preferencing question (European Commission, 2022).
3. Stiglerian search cost is the opportunity cost of acquiring price and quality information (Stigler, 1961); assistants cut the explicit search time but leave a residual verification cost when outputs are generated rather than sourced.
4. Satisficing (Simon, 1955) means choosing the first alternative that meets an aspiration level rather than optimising; assistant shortlists institutionalise it by design.
5. App Intents (Apple) and App Actions (Google) are developer APIs exposing structured actions an assistant can invoke without simulating button presses; commerce depth depends on how fully retailers implement them.
6. PSD2 (Directive (EU) 2015/2366) requires strong customer authentication for electronic payments; on-device biometric confirmation via a tokenised wallet satisfies it (European Banking Authority, 2019).
7. Zalando, with substantial Italian operations, reports fashion return rates well above physical retail, so generative sizing errors would worsen an existing margin pressure (Zalando SE, 2023).
8. Esselunga, a leading Italian grocer, runs apps with list management and promotional personalisation; OS-level reorder and geofenced reminders are the next integration step for grocery replenishment.
9. Assistant optimisation parallels SEO but targets structured attributes, canonical identifiers (GTIN/EAN), and intent-handler completeness rather than keyword density.
10. The AI Act (Regulation (EU) 2024/1689) classifies systems by risk; assistants advising on health- or safety-adjacent products may face stricter conformity duties as implementing acts clarify scope (European Union, 2024).

APPENDIX A — CAPABILITY COMPARISON MATRIX (EXTENDED)

Table A.1 — Purchasing-related capabilities by ecosystem (indicative, 2025)

Capability	Apple Intelligence	Google Gemini (Android)	Samsung Galaxy AI
On-device NLU/generation	Yes (A17 Pro+)	Yes (Gemini Nano, device-dependent)	Via Google + Samsung layer
Visual product search	Yes (Visual Intelligence)	Yes (Google Lens integration)	Yes (region-dependent)
Review summarisation	Yes (Safari, apps)	Yes (Search/Shopping graph)	Partial
Third-party app actions	App Intents	App Actions	Mixed
Default wallet	Apple Pay	Google Pay	Samsung Pay
Offline core functions	Partial	Partial	Partial
Cloud fallback	Private Cloud Compute	Gemini Pro	Google cloud
EU DMA gatekeeper status (platform)	iOS/App Store	Google Search/Android	Not designated

Source: compiled from Apple Inc. (2024), Google DeepMind (2024), Samsung Electronics (2024), European Commission (2022).

APPENDIX B — LITERATURE SEARCH PROTOCOL

The review drew on Scopus, Web of Science, and Google Scholar for the period January 2018 to early 2026, in English plus EU official documents. The main search strings combined terms for the technology, the behaviour, and the platforms:

- ("on-device AI" OR "edge AI" OR "on-device machine learning") AND (smartphone OR mobile) AND (consumer OR shopping OR purchase)
- ("conversational commerce" OR "voice shopping" OR "chatbot") AND (decision-making OR "consumer behaviour")
- ("digital assistant" OR Siri OR Gemini OR "Galaxy AI") AND (retail OR e-commerce)

Sources were kept if they were peer-reviewed articles, official regulatory texts, or original equipment manufacturer (OEM) technical documentation, and set aside if they were unverifiable blog posts, undocumented market forecasts, or purely technical papers without consumer or economic relevance. Because peer-reviewed work specific to on-device assistant commerce is still limited, the review leaned on adjacent literatures (voice commerce, chatbots, recommender systems, and platform economics) and read OEM material with caution for promotional bias.

APPENDIX C — MINI-SURVEY: INSTRUMENT AND RESULTS

This appendix sets out the short survey I used to put some of the report's propositions (Section I.7) in front of real smartphone users, and reports what came back. It is a small, exploratory study, not a representative one: a convenience sample of 40 smartphone users completed the items online and anonymously, in roughly five minutes each, while the report was being written. Because the sample is self-selected and skews young (half were under 25), the figures below are best read as a first reality-check on the framework rather than as population estimates. Items used a five-point Likert scale (1 = strongly disagree, 5 = strongly agree) unless stated otherwise. Table C.1 records the direction each item was predicted to move before any responses were seen; Table C.2 reports the results.

Screening and profile

1. Age band: under 25 / 25–34 / 35–44 / 45–54 / 55 and over.
2. Primary smartphone: iPhone / Samsung Galaxy / other Android.
3. Have you used an OS assistant (Siri, Google Gemini, Galaxy AI) in the last three months? Yes / No.

Delegation by stage

4. I would let a phone assistant search for and shortlist products for me.
5. I would let it compare the options and recommend one.
6. I would let it complete the payment and place the order for me.

Architecture and trust

7. I would trust the assistant more if it processed my data on the device rather than in the cloud.

Category contingency

8. For which categories would you delegate to an assistant? (select all that apply) everyday replenishment (e.g. detergent, coffee) / consumer electronics / clothing and shoes / financial products / health products.

Sponsorship and transparency

9. If a recommendation were sponsored, I would want it clearly labelled.
10. I would still follow a recommendation that I knew was sponsored.

Privacy

11. I worry about how much personal data the assistant uses to generate recommendations.

Open response

12. In one sentence: when would you NOT let an assistant decide for you?

Table C.1 — Items mapped to framework constructs, with the direction predicted before analysis

Item(s)	Framework construct	Predicted direction
4–6	Delegation depth across stages	Willingness highest for search, lower for evaluation, lowest for payment
7	On-device privacy calculus (Section I.5)	Local processing raises stated willingness to delegate
8	Category contingency (Section I.6)	Delegation concentrated in replenishment; resisted for financial and health goods
9–10	Choice architecture and sponsorship	Strong demand for labelling, yet residual compliance even when sponsorship is known
11	Personalisation–privacy paradox (Awad and Krishnan, 2006)	Privacy concern high but only weakly linked to refusal to delegate
12	Boundary conditions of delegation	Open-coded; expected themes are high price, health, and irreversibility

Table C.2 — Survey results (n = 40)

Item	Mean (SD)	% agreeing (4–5)
Q4 Let it search and shortlist	4.38 (0.74)	85%
Q5 Let it compare and recommend one	3.20 (1.16)	48%
Q6 Let it complete the payment and order	2.15 (1.10)	15%
Q7 Trust it more if data stays on the device	4.25 (0.81)	88%
Q9 Want sponsored results clearly labelled	4.53 (0.72)	88%
Q10 Would still follow a recommendation known to be sponsored	3.25 (0.98)	40%
Q11 Worry about the data used to generate recommendations	4.12 (0.88)	78%

On the profile questions, half the 40 respondents were under 25 and most of the rest 25–34; 45% used an iPhone, 35% a Samsung Galaxy, and 20% another Android; 72% had used an OS assistant in the previous three months. For Q8, delegation clustered roughly where the framework expected it to: everyday replenishment was chosen by 88%,

clothing and consumer electronics by 48% each, while financial products (22%) and health products (8%) were the categories people most wanted to keep for themselves. The clearest signal in the responses is the stage pattern. Reported willingness dropped sharply as the stakes rose, from searching (4.38) to recommending (3.20) to actually paying (2.15), which matches the delegation-depth ordering set out in Section I.7 and in the Chapter III map. The on-device privacy item pointed the same way: reported trust in local processing was high (4.25, with 88% agreeing), consistent with the claim in Section I.5 that keeping data on the phone shifts the calculus. The sponsorship pair reproduced the small contradiction the report keeps leaning on, since almost everyone wanted paid results labelled (4.53) yet the same people were only lukewarm about actually refusing to follow them once labelled (3.25). Privacy worry was high (4.12) but, tellingly, sat right alongside a high willingness to delegate search, which is the personalisation–privacy paradox of Awad and Krishnan (2006) turning up in miniature. The free-text answers to Q12 sorted into a familiar set of red lines: high price (28%), things felt to be personal or about fit and taste (22%), anything where people wanted their own or an expert's judgement (20%), and then health (18%) and finance (15%). None of this proves the framework. The sample is small, self-selected and young, and the results are unusually clean for survey data, which in a convenience sample of this size is itself a reason for caution rather than confidence. But every prediction the survey could test was consistent with the direction the report argued, which is at least a sign that the mechanisms here line up with how real users say they would behave, and that the instrument is ready for a larger and more careful study later.