

Beyond the Numbers

Personal AI Agents and Disruptive Narratives

Thomas Shipp, CFA
Head of Equity Research

Tucker Beale
Sr. Equity Research Analyst

Executive Summary

Meta's September launch of Muse, a personal artificial intelligence (AI) agent designed to act on a user's behalf across digital workflows, crystallized the market's latest disruption narrative. If AI can make consumer switching easier, which business models are truly protected by moats and which are merely monetizing friction? Predicting consumer behavior is generally not where we find our edge, but when the market believes a new technology will break down economic moats, we pay attention. While we are skeptical that AI agents will disrupt businesses whose moats were earned through differentiated products, services, or brands, we do see disruption risk for businesses that largely benefit from consumer inertia, like forgetting to cancel subscriptions or not wanting to spend time on price comparisons. In this month's *Beyond the Numbers*, we unpack how personal AI agents could disrupt online consumer-facing businesses and analyze the business model of personal AI agent software platforms.

Key Takeaways

- **The personal AI agent disruption risk is to consumer inertia, not to earned moats.** Businesses protected by strong products, services, or brands look less exposed than those that profit from consumers not shopping around, such as unused subscriptions, low-rate deposits, and passive insurance renewals.
- **Expect partial disruption, not an extinction event.** Personal AI agents are more likely to take a share of discovery and comparison than to replace online travel agencies, banks, or e-commerce platforms outright.
- **The "SaaSocalypse" shows that narratives can overshoot fundamentals.** Software stocks sold off on AI disruption fears even as software sales growth sped up to ~17–18% year over year. Investors should separate market imagination from financial evidence, because the income statement will show whether personal AI agents are actually changing behavior.
- **The economics of personal AI agents are unproven.** Only 2–3% of consumers currently pay for AI. Even a very bullish one billion subscriptions would bring in less than half of the estimated 2026 AI capex. With possibly weak switching costs and little brand differentiation, these tools may end up as features built into operating systems rather than standalone profit engines.

Help Wanted: Personal Assistant, Humans Need Not Apply

The September 8 release of Meta Platforms' (META) personal AI agent Muse brought long-hyped, but until recently hypothetical, personal AI agents into the forefront of investor's minds. But what is a "personal AI agent"? Broadly speaking, it is a large language model (LLM) powered software platform that can carry out long-running tasks on a user's behalf rather than simply answering questions in a chat window. A chatbot like ChatGPT tells you which flights are available with which carrier, and which is the cheapest, while a personal AI agent like Muse can search, compare, fill out forms, and book the flight on your behalf, and keep working after you close the app. A task such as this is not ground-breaking on its own, but the disruption comes from bundling many of these sorts of tasks into a persistent online personal assistant that never gets tired of the tedium of modern life.

The current generation of consumer-focused personal AI agents is early, and probably still overhyped, but the strategic direction is clear; the largest technology platforms are aiming to become the default interface through which consumers manage online economic decisions.


Of the personal AI assistant/agent products, Muse is the most visible example. OpenAI's Dots, Google's Gemini Spark, xAI's Grok Bot, and other consumer-facing agent products are all variations on the same idea. In the rapidly evolving market for AI-driven tools, the consumer-facing focus has seemingly shifted overnight from building models that give better answers to monetizing a user's desire to delegate.

Now is a good time to differentiate between the consumer-focused "personal AI agent" class of AI agents, which is the focus of this report, and enterprise-focused productivity and workflow "agentic AI" tools, such as Anthropic's Claude Cowork alongside any number of proprietary, homegrown efforts. To generalize, enterprise AI agents focus largely on computer-based "knowledge work" tasks, such as the following prompt: "find the August Board PowerPoint presentation, update titles, date references, data and analysis with September data, save the draft as a PDF, and send to my team via email". Businesses around the world are working to integrate these agentic AI tools across their organizations to enhance worker productivity, and that potential productivity is more or less the promise of AI at the enterprise level.

Potential Disruption from Your Personal AI Agent

Many consumer businesses earn attractive returns because customers do not revisit their purchasing decisions very often. This kind of consumer inaction is monetized across diverse business models throughout the economy, such as low-rate bank deposits, unused subscriptions, insurance renewals, and price-comparison businesses. Goldman Sachs has created a thematic "Consumer Inertia" stock index that is built around exactly this idea. It contains U.S.-listed companies in industries where retention is supported by behavioral, operational, or financial switching costs. AI disruption narratives, correct or incorrect, take hold quickly in today's market. "Consumer Inertia stocks have lagged since the release of Muse" charts equity market performance since the end of August (five trading days before the release of META's Muse) through October 9. During this brief period, the "Consumer Inertia" index of stocks has shed 11.5%, compared to the S&P 500's gain of 1.6%, and Goldman's "Agentic AI" thematic index gain of 9.7%.

Consumer Inertia Beneficiaries Have Lagged Since the Release of META's Muse



Source: LPL Research, Goldman Sachs, Bloomberg 10/09/26

All indexes are unmanaged and cannot be invested in directly. Past performance is no guarantee of future results.

[Joint research](#) from economists at Stanford University and Texas A&M University studying subscription behavior found that when card replacement forced customers to make an active renewal decision, the monthly drop in

retention was four times larger than in passive renewal periods. Their research concluded that consumer inertia raised seller revenue by 87% to 120% on average. If agents make active renewal and active comparison the default, they could pressure the inertia embedded in many business models.

E-commerce

First up, e-commerce, which is among the cleanest examples because the current model depends heavily on discovery as consumers search, compare, and eventually make a purchasing decision. In a zero-click agentic commerce model, the consumer may never visit a retailer's website, never see a search-results page, and never encounter the sponsored listing that currently funds much of the discovery ecosystem. Bullish industry analysts have described this as a 2026 shift in which shoppers increasingly use generative AI (i.e., LLMs) channels to find and buy products, and eventually delegate the transaction itself to autonomous agents.

Potential disruption is not uniform, however. Strong brands, low-price leaders, and retailers with proprietary inventory will likely continue to do well. The aggregation layer, such as lead-generators, marketplaces, and ad-funded discovery platforms that monetize the consumer's search process, is more vulnerable. Morgan Stanley has projected that nearly half of online shoppers could use AI shopping agents by 2030, representing approximately 25% of their spending. That estimate may prove too aggressive, but they explain the market's concern. Our view is that LLMs can certainly add value to the discovery process (finding deals, parsing through reviews), which can drive some level of ad revenue disruption, similar to broader web search economics. However, we recommend a healthy dose of caution on extrapolating "LLM-enhanced search and discovery" to "agentic commerce becomes the primary way consumers transact online".

Online Travel Agencies

Online travel agencies (OTAs) are cited in nearly every personal AI agent disruption scenario — alongside e-commerce — and is worth scrutiny. OTAs such as Expedia and Booking Holdings are essentially search, aggregation, and booking convenience layered over fragmented supply of flights, hotel rooms, experiences, and rental cars. Why search and price-compare on multiple OTA websites, when your agent can do that for you? Therefore, we understand why shares of Expedia and Booking Holdings sold off more than 3.5% when META reported that Muse was downloaded 2.5 million times in two weeks.

The OTA bear case suggests that if an agent can search across airlines, hotels, loyalty balances, and payment credentials, it can reduce the consumer's need to visit one or multiple travel marketplaces. Travel is a high-cost, exception-heavy, and at times emotional, purchase. People care about loyalty points, seat assignments, room/property locations, and reviews. Further, people want assurances around cancellation and refund policies, and need support when something goes wrong. The [June 2026 traveler survey](#) from global consultancy Accenture found that 87% of consumers were open to working with AI agents on travel decisions, but only 10% were willing to let an AI agent handle the actual payment without human approval. This leads us to the conclusion that personal AI agents may disrupt some level of travel discovery rather than fully displace OTA platforms.

Financial Institutions: Bank Deposits and Insurance Policies

Traditional deposit franchises have long benefited from customers who leave cash in low-yield accounts even when higher-yield alternatives might be available. The most likely reason for this inaction, in our view, is the hassle of changing banks and the trust built with existing institutions, not that consumers don't know that there are higher cash yields available. A personal AI agent claims to solve the hassle friction directly, by simply doing it for you. However, we're not so sure consumers are ready to place similar levels of trust and comfort with big tech companies. Our view is that the personal AI agent risk to banks is marginal and financial, not disruption. The earnings risk is that deposits become more rate-sensitive, driving greater competition for those deposits (via higher yields).


Insurance policy renewal shopping is another area where agents could reduce the value of consumer inertia. Auto, homeowners, renters, and small-business insurance often renew with limited active comparison, partly because the process is tedious. A personal AI agent could turn that annual renewal from a passive rollover into an active auction. The exposed economics would sit most directly in insurance lead generation and comparison marketplaces, but carriers with high retention supported more by customer inertia than by differentiated underwriting, service, or bundling could also face higher shopping rates and more price-sensitive renewal behavior.

Parallels to the “SaaSpocalypse”

We don’t have to look back very far to find a useful parallel to the latest AI disruption stock selloff. In early February 2026, software stocks sold off sharply on the fear that AI would collapse large parts of the “Software-as-a-Service” (SaaS) market, an event now widely known as the “SaaSpocalypse.” The narrative was AI coding tools could create all the software an enterprise needs and remove the need for seat-based software subscriptions, leaving only the strongest survivors. The fundamentals did not validate the panic. For one example, Salesforce beat and raised full-year revenue guidance in their recent two-fiscal quarter earnings reports and issued long term growth targets well above consensus estimates. This continued to play out sector-wide, as many software companies came in ahead of consensus revenue estimates — with Palantir, Snowflake, and Datadog — beating by 4–7% over the last two quarters, while SaaS revenue (measured using the trailing 12 month sales of the S&P North American Expanded Technology Software Index) grew ~17–18% year over year, the fastest pace in several years. Analyst earnings estimates for fiscal 2027 followed the narrative, declining about 10% from the beginning of the year to late March, going from a 10% estimated year-over-year growth rate down to -5%. From there, 2027 EPS estimates have effectively retraced the entire decline, now sitting approximately at the same level as at the start of the year.

A disruption narrative can make sense, and stock performance can overshoot the actual fundamental impact. The release of personal agents has legitimately pressured sentiment in the segments we discussed, but sentiment does not always map to reality. The right response is not to dismiss disruption risk, but to separate market imagination from financial evidence. The dual panel exhibit “SaaSpocalypse Narrative Hammered Stock Prices Despite Strengthening Fundamentals” charts the software industry’s stock underperformance relative to the S&P 500 index and Goldman’s “US Broad AI” thematic index in the first panel, and the software industry’s accelerating sales growth in the second panel.

“SaaSpocalypse” Narrative Hammered Stock Prices Despite Strengthening Fundamentals



Sources: LPL Research, Bloomberg, Goldman Sachs, 10/09/26

SaaS industry performance and sales growth represented by the S&P North American Expanded Technology Software Index fundamentals
All indexes are unmanaged and cannot be invested in directly. Past performance is no guarantee of future results.

Personal Agent Economics

We've covered several business models that personal AI agents could potentially disrupt, but what of the personal AI agent business model itself? Today, the easiest revenue models for personal AI agents are subscription sales, a usage-based model, or a combination of the two. That makes sense for early adoption, but it may not be the end state, especially given the low penetration of paid consumer AI subscriptions among U.S. households. The more powerful monetization opportunity is based on the transactions the agent undertakes on a user's behalf. If an agent books a flight or routes a purchase, it will eventually try to capture some of the economics currently earned by OTAs, marketplaces, and payment networks.

The problem, as usual, is cost and unclear unit economics at the AI model layer. Every long-running task can involve massive amounts of compute and other operating costs, including model inference, browser sessions (each user is assigned a dedicated virtual computer to accomplish this), app integrations, identity verification, and customer support. The economics are questionable even if you assume expanded penetration of users that actually pay for a flat-rate subscription. Estimates suggest only 2-3% of consumers currently pay for AI products, with average monthly spend around \$31. One billion personal AI agent subscriptions at \$31, an extremely bullish hypothetical estimate, implies annual revenue of ~\$372 billion. This would not cover even half of the estimated AI capex spend in 2026, much less the even greater expected spend in 2027 and beyond. The industry still has to prove that consumer willingness to pay for AI can support the cost of autonomous execution.

That cost structure helps explain why advertising, referral fees, and transaction takes will be tempting. This creates a second problem, in our view, and that is trust. If a consumer asks an agent to find the best checking account, the cheapest flight, or the right product, the agent's value proposition depends on perceived alignment. The moment recommendations are visibly sponsored, the product risks becoming a search with a conversational wrapper.

Personal AI Agent Platform Durability

There are paths to overcoming the unit economics concern. The personal AI agent market could evolve into a "winner-takes-most" platform with minimal competition. A first-mover advantage could accrue to a free offering (with expanded features available via paid subscriptions), providing a large user-base to "lock-in" and clearing a path to victory in the battle for market share. The first-order advantage and bull case for becoming a "winner-takes-most" platform is switching costs: as an agent accumulates preferences, permissions, workflow history, and trust, it becomes harder for the user to replace. The second-order advantage is scale: if one agent platform controls enough high-intent demand, suppliers may compete for access, creating a demand-aggregation flywheel that looks network-effect-like.

The bear case against a particular personal AI agent becoming a "winner-take-most" platform is that these tools in fact have *weak* switching costs. A competitor could offer a one-click import tool that transfers user configuration, history, and permissions into a new assistant. If that happens, agents may struggle to earn durable excess returns and instead become a thin layer of convenience on top of existing platforms. That would still matter for incumbents that depend on aggregation and forgetfulness / inertia, but it would make the agent business itself less attractive than the market currently projects. Ironically, should a personal AI agent platform gain mass consumer traction but exhibit weak switching costs, it will likely benefit from the exact same inertia it purports to exploit on behalf of consumers.

Our View

Our view is... nuanced. We don't see personal AI agents causing some sort of extinction event for e-commerce, OTAs, banks, or subscription businesses, at least not anytime soon. Over time, and if adoption reaches a large enough share of internet traffic, AI agents do present profit margin risks for businesses that rely on (1) aggregation and discovery, i.e., users navigating the internet (ad revenue requires eyeballs), and/or (2) consumer inertia, whether it be subscription revenue or financial products. Businesses with trusted brands, strong loyalty programs, or genuine service quality will likely prove more resilient than the early narrative suggests. Furthermore, similar to the

“SaaSocolypse” episode, we doubt the potentially disrupted businesses stand pat and hope the agents don’t disrupt them. These businesses will evolve, and some are likely developing agentic tools themselves. Those that don’t deserve the disruption.

As for the business of personal AI agents themselves, we’re not convinced there is anything in these tools that moves the needle in the monetization of consumer AI products. The same commodification issues present in monetizing consumer usage of chat-based LLM products do not go away with agent-based LLM personal assistants. Those issues — generally speaking — are these tools all do the same thing, they all perform similarly fine, and they have minimal brand power. Nor are we convinced that personal AI agents, with today’s capabilities, are ready for prime time. If a user does decide they need expanded features and upgrades to a paid personal AI agent, how many times do they forgive the agent that couldn’t perform a task before they cancel and go back to the free version? The AI industry is hoping personal AI agents are the “killer app” that can profitably break open the consumer AI market. We are not convinced. We see greater likelihood that the functionality provided by these tools ultimately ends up as features natively bundled into operating systems (whether mobile or desktop), powered by a multi-model routing system. However, when it comes to consumer business disruption due to personal AI agents, the fundamentals will decide whether the story is real.

Important Disclosures

This material is for general information only and is not intended to provide specific advice or recommendations for any individual. There is no assurance that the views or strategies discussed are suitable for all investors or will yield positive outcomes. Investing involves risks including possible loss of principal. Any economic forecasts set forth may not develop as predicted and are subject to change.

References to markets, asset classes, and sectors are generally regarding the corresponding market index. Indexes are unmanaged statistical composites and cannot be invested into directly. Index performance is not indicative of the performance of any investment and do not reflect fees, expenses, or sales charges. All performance referenced is historical and is no guarantee of future results.

Any company names noted herein are for educational purposes only and not an indication of trading intent or a solicitation of their products or services.

All information is believed to be from reliable sources; however, LPL Financial makes no representation as to its completeness or accuracy.

Because of their narrow focus, sector investing will be subject to greater volatility than investing more broadly across many sectors and companies.

Earnings per share (EPS) is the portion of a company’s profit allocated to each outstanding share of common stock. EPS serves as an indicator of a company’s profitability. Earnings per share is generally considered to be the single most important variable in determining a share’s price. It is also a major component used to calculate the price-to-earnings valuation ratio.

Goldman Sachs’ “consumer inertia” index is a thematic equity portfolio of companies that profit from customers’ reluctance to switch providers or cancel services.

The S&P North American Technology Software Index is a sector-specific equity index designed to measure the performance of publicly traded software companies in North America.

All index data from Bloomberg.

This Research material has been prepared by LPL Financial LLC.

Not Insured by FDIC/NCUA or Any Other Government Agency	Not Bank/Credit Union Guaranteed	Not Bank/Credit Union Deposits or Obligations	May Lose Value
---	----------------------------------	---	----------------

Tracking #1190724 | #1190729 (Exp. 10/2027)