

Competition in Generative AI: Evidence from Consumer Demand^{*}

Mark Egan,[†] Dominic Russel,[‡] Claire Shi,[§] and Vincent Yao[¶]

October 16, 2026

Abstract

How competitive will the market for generative AI be as it matures? Using novel transaction-level credit and debit card data covering 174 million consumers and 114 billion transactions, we study the adoption of paid services and the demand-side foundations of competition in this market. As of April 2026, approximately 2.9% of U.S. consumers pay for at least one major generative AI service, with adoption highest among younger and higher-income consumers. Nearly all of these paid users subscribe to basic plans that cost \$30 or less per month. Three patterns suggest limited differentiation and low switching costs. First, although multihoming is common, it has declined as providers' capabilities have expanded. Second, annual customer retention rates range from 20% to 45% across major AI providers, lower than those for many other subscription products. Third, switching rises around major model releases. Motivated by these findings, we develop and estimate a structural model of AI service demand and supply that reveals strong competitive pressure. Consumers are highly price sensitive, have much weaker preferences for particular providers than for AI overall, and view AI services as increasingly close substitutes. Our results suggest that, absent a substantial increase in AI model quality differentiation, growing substitutability among services will constrain providers' pricing power and allow consumers to capture a large share of the surplus generated by these services.

^{*}First version: October 2026. We thank Toren Fronsdal, Sam Hanson, Richard Ruback, Jesse Shapiro, and Adi Sunderam, and seminar and conference participants at Harvard Business School for helpful comments. Vincent Yao is an unpaid member of the Fiserv Advisory Council and a paid consultant. Mark Egan is an unpaid member of the Fiserv Advisory Council.

[†]Harvard Business School & NBER. Email: megan@hbs.edu.

[‡]Harvard Business School and Department of Economics. Email: drussel@g.harvard.edu.

[§]Harvard Business School and Department of Economics. Email: cshi@g.harvard.edu.

[¶]J. Mack Robinson College of Business at Georgia State University. Email: wyao2@gsu.edu.

1 Introduction

Recent advances in generative artificial intelligence have the potential to transform economic activity and daily life. Since the release of OpenAI’s ChatGPT in late 2022 and Anthropic’s Claude in 2023, generative AI has rapidly become a prominent consumer technology used for a wide variety of everyday tasks, ranging from writing and coding to research and education. The adoption and pricing of these services, their effects on consumer welfare, and the valuations of AI firms will depend on how competitive the market becomes.

A central question is therefore whether providers of consumer-facing generative AI services can sustain market power as the market matures. The answer depends in part on consumer demand, particularly on how differentiated consumers perceive services to be and how readily they substitute one service for another. Yet, despite generative AI’s rapid growth and economic importance, systematic evidence on consumer demand and competition remains limited. In particular, little is known about who pays for these services, how consumers choose among providers, how persistent paid subscriptions are, whether services are substitutes or complements, and how responsive consumers are to changes in price and quality.

This paper uses novel, de-identified transaction-level credit and debit card data to study consumer adoption of and demand for paid generative AI services. We first document how adoption varies across demographic groups and regions and how it has evolved over time. We then examine purchasing dynamics, including how long paid subscriptions persist, what leads consumers to switch providers, and whether they multihome. These patterns shed light on the degree of differentiation among services and consumers’ willingness to substitute across them—two demand-side forces that are critical to understanding competition, consumer welfare, and the market’s long-run structure.

Motivated by these empirical patterns, we develop and estimate a structural model of AI service demand and supply. The demand side of the model builds on [Gentzkow \(2007\)](#) and uses the panel structure of our data to separately identify consumers’ preferences for individual services, the correlation in those preferences across services, and the incremental utility from subscribing to multiple services. Distinguishing correlated preferences from complementarity is important. A consumer may subscribe to both OpenAI and Anthropic products because of a strong general taste for AI, or because they are complements so that subscribing to one raises the value of the other. These features of demand have distinct implications for competition. We use the estimated model to quantify how these forces shape market power, measure the surplus generated by AI services and its division between consumers and providers, and examine how changes in market structure affect competition.

Overall, our reduced-form and structural results point to strong demand-side competitive pressure in the market for consumer AI services. Consumers differ substantially in their overall taste for AI, but provider-specific tastes are small relative to this common component: for

instance, consumers who value an OpenAI product also tend to value an Anthropic product. Moreover, as model capabilities have expanded and converged, consumers increasingly treat these services as substitutes. Access to one service is therefore less valuable to a consumer who already has access to another. Together with consumers' price sensitivity, these patterns leave providers with limited scope to sustain market power through provider-specific preferences. Although consumers' sensitivity to model quality could give a provider market power if it achieves a level of quality that rivals cannot match, frontier model capabilities have thus far converged across providers. Absent a reversal of this pattern, our results suggest that the demand side is unlikely to support sustained market power and that, ultimately, consumers will capture a large share of the surplus generated by generative AI products.

Our transaction-level sample spans January 2021 through April 2026 and contains 174 million consumers and 114 billion transactions. Transaction descriptors allow us to identify payments to major AI providers. We focus primarily on the four large generative AI providers observed in our data: Anthropic (Claude), Google (Gemini), OpenAI (ChatGPT), and xAI (Grok). For each transaction, we observe the date and amount, together with anonymized identifiers for the card and consumer. These data allow us to construct a panel that tracks consumers' purchases over time and across AI providers. We also observe demographic information, including consumers' ages, incomes, and regions.

We focus on paid consumer adoption, excluding both free AI subscriptions and enterprise adoption. Unlike free subscriptions, paid AI subscriptions reflect providers' ability to monetize consumer demand, and are a core source of their revenue. For OpenAI, consumer subscriptions account for approximately 75% of revenue in 2024 and 60% at the beginning of 2026. Even Anthropic, which has traditionally targeted enterprise customers, generates approximately 20% of its revenue from consumer sales.¹ Consumer subscriptions may also be especially important for future improvements in model quality, since AI companies generally do not train models on enterprise data.

We use the data to document several new facts about consumer adoption of paid AI services. First, paid adoption has grown but remains relatively limited in the United States. In April 2026, 2.9% of consumers paid for at least one product from one of the four major AI providers in our sample, in line with recent industry estimates.² OpenAI was the most widely adopted provider, with 2.4% of consumers paying for an OpenAI product. Although the share of consumers paying for Anthropic products had grown in recent months, it remained at only 0.5% in April 2026. Paid adoption is substantially more common among younger and higher-income consumers: those earning more than \$200,000 per year are 1.7 times as likely to pay as those earning less than \$50,000, and those younger than 55 are 1.8 times as likely as those aged 55 or older. These relationships between demographics and adoption hold broadly across AI providers. We also

¹See estimates from [Gavin \(2026\)](#); [Parham \(2026\)](#); [Lin \(2025\)](#).

²See, for example, [Bank of America Institute \(2026\)](#); [PNC Economics Research \(2026\)](#).

document meaningful geographic variation, with adoption highest in California, Washington, and the Mid-Atlantic region.

Second, consumers frequently multihome across AI providers, particularly during the earlier part of our sample. For example, roughly 45% of consumers who paid for an Anthropic product also paid for an OpenAI product in June 2024. However, the prevalence of multihoming has declined over time. By April 2026, less than 30% of consumers who paid for an Anthropic product also paid for an OpenAI product. This decline is consistent with AI services becoming less complementary as providers have developed increasingly general-purpose products with overlapping capabilities.

Third, retention among paying customers is low, especially relative to other subscription products such as newspapers and streaming services. We measure retention by following a consumer’s first observed purchase from an AI provider. Retention varies across providers and is highest for OpenAI: after 12 months, roughly 45% of OpenAI users remain subscribed, compared with 30% of Anthropic users. By comparison, annual retention rates for Netflix and Spotify in our data are around 70%.

Fourth, consumers switch across providers, and switching activity increases following major model releases. We measure model capabilities using data from the Epoch Capabilities Index ([Epoch AI, 2025](#)). When Anthropic or OpenAI releases an improved model, the share of consumers who paid only for the rival in the previous month and begin paying for the provider with the model improvement rises by about 25%.

Taken together, these facts describe a growing market characterized by limited customer loyalty and services that appear increasingly substitutable. If these patterns persist, they suggest that the market for AI services will likely be highly competitive and it will be challenging for providers to sustain market power in the long run.

Motivated by these reduced-form facts, we develop and estimate a structural model of demand for AI services to better understand consumer preferences, how those preferences have evolved over time, the resulting welfare consequences, and the implications for market structure. Following [Gentzkow \(2007\)](#), we model consumer choice as a discrete-choice problem in which consumers choose among bundles of AI services rather than necessarily selecting a single service. The model allows consumers to purchase multiple services simultaneously and permits pairs of services to be either complements or substitutes.

We allow consumers to have heterogeneous and correlated preferences across AI services in the spirit of [Berry *et al.* \(1995\)](#) and [Berry *et al.* \(2004\)](#). Distinguishing preference heterogeneity from complementarity is particularly important given that either force can explain the prevalence of multihoming. For example, a consumer who places a high value on Anthropic’s Claude may also place a high value on OpenAI’s ChatGPT because the consumer has a strong general preference for AI services, even if the two services are substitutes. Alternatively, the services may be complements, such that subscribing to Claude increases the value of subscribing

to ChatGPT.

On the supply side, we assume that AI providers compete in a differentiated-products Bertrand–Nash price-setting game. Providers recognize that consumers may purchase bundles containing products from multiple firms and that the price of one product can therefore affect demand for other products. When products are substitutes, a price increase by one provider shifts demand toward its competitors. When products are complements, a price increase can reduce demand for both products. Thus, greater substitutability intensifies price competition, whereas greater complementarity softens competition and supports higher equilibrium prices.

We also consider a dynamic extension with learning by doing on the supply side. Specifically, we allow future model capabilities to depend on the number of users a provider attracts today, reflecting the idea that a larger user base gives the provider access to more data or resources that can be used to improve model quality (Prufer and Schottmüller, 2021; Hagiu and Wright, 2023). This force creates a dynamic incentive to set lower prices today: lower prices attract more users, improving future model quality and raising future demand.

We estimate the demand model by maximum likelihood using our transaction-level panel data on consumer purchase decisions. The panel structure is central to separately identifying persistent heterogeneity in consumer preferences, correlations in preferences across providers, and complementarity or substitutability between AI services. Repeated purchase decisions and switching patterns identify persistent differences in consumers’ preferences. Correlated purchase behavior across providers helps identify correlations in preferences, while, conditional on those preferences, consumers’ tendency to multihome rather than switch between services helps identify whether the services are complements or substitutes.

Our structural estimates on the demand side provide insight into the nature of consumer preferences and how they have evolved over time. To begin, consumer demand for AI services is relatively price elastic, with an average own-price elasticity of about -4.6. In our model, this high elasticity implies relatively little horizontal differentiation across AI providers and thus limited market power for providers.

Our estimates of consumer preference heterogeneity reinforce the conclusion of limited horizontal differentiation. Although preferences vary substantially across consumers, consumer-specific preferences are strongly correlated across AI providers: consumers who value one provider more than average also tend to highly value other providers. For example, the estimated correlation between consumer-specific preferences for OpenAI and Anthropic products is 0.89. The lowest preference correlation between any pair of products is OpenAI and Google, at 0.55, which is still substantial. These estimates suggest that providers largely appeal to the same consumers and that provider-specific preferences account for relatively little of the variation in demand.

Despite this limited horizontal differentiation, we find some degree of vertical differentiation across providers, consistent with observed market shares. At the end of our sample, the

average consumer was willing to pay a \$10 monthly premium for OpenAI products relative to Anthropic products. Part of this difference reflects model quality: our estimates indicate that consumers are willing to pay more for higher-quality models. For instance, moving from GPT-4’s capability level to GPT-5.5 Pro’s raises average willingness to pay for OpenAI products by approximately \$25 per month.

Finally, we generally find that consumers increasingly view AI services as substitutes, particularly OpenAI and Anthropic products. As the capabilities of these services have expanded and converged, the incremental value of subscribing to one service for a consumer who already subscribes to the other has declined. This evolution is consistent with the observed decline in multihoming and has important implications for competition and the long-run structure of the consumer AI market. In sum, our demand estimates point to three sources of competitive pressure: limited horizontal differentiation, strong substitution in response to quality improvements, and the declining incremental value of additional subscriptions. Under differentiated-products Bertrand–Nash pricing, these forces put downward pressure on equilibrium markups.

We then use our demand estimates to calibrate the supply side of the model. We estimate markups, expressed as a share of the price (i.e., Lerner Index) ranging from 19% to 59% across products, suggesting that about 41-81% of revenues are used to cover the marginal costs of running these AI models. In our dynamic extension with endogenous improvements in model quality, we find even lower markups. In particular, when providers internalize the future value of expanding their user bases, the marginal cost of operating a typical service is roughly 70% to 80% of its subscription price. These estimates capture the variable marginal costs of providing AI services but exclude fixed operating costs and expenditures on research and model development. Thus, they are an upper bound on total profitability.

The model also allows us to conduct counterfactual analyses and quantify the consumer surplus generated by paid AI services. We estimate that access to paid AI subscriptions generates average consumer surplus of approximately \$1.09 per consumer per quarter. This average masks substantial heterogeneity because most consumers do not currently subscribe to an AI service. Among adopters, paid AI services generate average consumer surplus of approximately \$35 each quarter, over and above the amount spent on AI subscriptions.

We also use the model to examine how consumer preferences for AI services shape competition. We begin by considering substitutability across services. By the end of our sample, all pairs of AI services are substitutes on average: the marginal utility of an additional subscription is lower for a consumer who already subscribes to another service. For some services, such as Anthropic which is very substitutable with OpenAI, equilibrium prices would be 30% higher under a counterfactual with no substitutability, and up to 60% higher if the services were complements, as several service pairs were early in our sample.

The strong correlation in consumers’ preferences across providers also constrains markups by limiting product differentiation. In a counterfactual in which preferences across providers

are uncorrelated, and services are therefore more differentiated, equilibrium prices are approximately 2% to 10% higher.

Finally, we examine Anthropic’s effect on competition. Although OpenAI remains the dominant provider, Anthropic has grown rapidly from a small base since the release of Claude Code. To quantify Anthropic’s competitive role, we consider a counterfactual in which Anthropic is removed from the market. Because consumers’ preferences for the two providers’ services are highly correlated and the services are strong substitutes, Anthropic exerts meaningful competitive discipline on OpenAI. In the counterfactual, OpenAI would raise its prices modestly, but sufficiently that its equilibrium market share would fall relative to its market share in the baseline scenario when it competes with Anthropic. The joint price and adoption response is consistent with Anthropic constraining OpenAI’s exercise of market power. Accounting for the fact that other providers change prices in response to entry amplifies the welfare gain of having Anthropic in the market by 21%, to \$5.05 per quarter per AI user.

Taken together, our results provide new evidence on consumer demand for generative AI services. Despite AI’s prominence in public discourse, paid consumer adoption remains modest, although recent trends suggest that it will continue to grow rapidly. The limited persistence of paid subscriptions, correlation in preferences, and the shift from complementarity toward substitutability have important implications for the diffusion of AI services and the industry’s long-run competitive structure.

Related Literature

Our paper contributes to a rapidly growing literature on the adoption and use of generative AI. Recent evidence suggests that exposure to generative AI is now common among both consumers and workers. Survey estimates (e.g., [Bick et al., 2026a](#); [Gottfried et al., 2026a](#)) show that around half of U.S. adults have used generative AI, and that around a quarter to half of U.S. employees have used generative AI for work (e.g., [Bick et al., 2026c](#); [Board of Governors of the Federal Reserve System, 2026](#); [Kemp, 2025](#); [Bick et al., 2026b](#)). AI adoption for business tasks is also growing inside the firm ([Bonney et al., 2026](#)), and at the executive level ([Yotzov et al., 2026](#)). Our analysis complements these survey-based studies by using transaction-level panel data to identify payments to specific AI providers and follow the same consumers over time. This allows us to study provider-specific adoption, multihoming, switching, and retention, and their implications for competition. Our measure is also deliberately narrower than survey-based measures which capture free and enterprise use.

Others use non-survey data to study AI adoption. Among these, descriptive industry reports ([Bank of America Institute, 2026](#); [PNC Economics Research, 2026](#)) are the closest to our use of transaction-level data. Using aggregated Bank of America data, [Bank of America Institute \(2026\)](#) reports that roughly 3% of households pay for AI services and that paid adoption is concentrated among younger and higher-income households, consistent with our first descrip-

tive findings. [Demirer et al. \(2025\)](#) use API usage data from OpenRouter and Microsoft Azure (platforms primarily serving developers and technology firms) to document rapid entry, falling prices, and frequent turnover among leading LLMs. We find similar evidence of substantial competitive pressure in the consumer AI market. [Caoui \(2026\)](#) uses data from LMArena, where users choose between anonymized LLM responses, to study preferences over model capabilities and response styles. Data from AI providers have also been used to analyze how consumers use AI products: for example, [Chatterji et al. \(2025\)](#) and [Handa et al. \(2025\)](#) find that writing, information, and software development are particularly common use cases.

Second, we contribute to emerging work quantifying the impacts and value of generative AI. Most closely related is [Brynjolfsson et al. \(2026\)](#) who use online choice experiments to elicit the compensation consumers would require to give up access to generative AI chatbot tools for one month. In March 2026, they estimate a mean willingness to accept of \$124.50 per month and a median willingness to accept of \$11.40. Their welfare object is complementary, but different from ours. Their estimates capture the value of access to generative AI tools, including services available without a paid subscription, whereas we estimate the incremental surplus generated by paid products relative to an outside option that may include free or employer-provided access. More broadly, we relate to a number of recent works that study the effects of AI tools on economic outcomes such as productivity and firm performance (e.g., [Noy and Zhang, 2023](#); [Chen and Chan, 2024](#); [Cui et al., 2024](#); [Chen and Stratton, 2026](#); [Otis et al., 2024](#); [Dillon et al., 2025](#); [Song et al., 2026](#); [Sarkar, 2025](#); [Brynjolfsson et al., 2025](#); [Hampole et al., 2025](#); [Babina et al., 2026](#)), as well as the impact of consumers delegating purchase choices (e.g., [Delaprez and Hortaçsu, 2026](#)) and financial decision-making (e.g., [Choukhmane et al., 2026](#); [Lu et al., 2025](#)). Building on this literature, we study the welfare implications of AI market structure, complementing existing work on the potential welfare impact of increased productivity, worker displacement, or changes in economic behavior (e.g., [Athey and Morton, 2025](#)).

Finally, our paper relates to the industrial organization literature on demand estimation, new products, and consumer welfare ([Berry et al., 1995, 2004](#); [Nevo, 2000](#)). [Petrin \(2002\)](#) demonstrates how demand-estimation methods and consumer-level information can be used to quantify the welfare gains generated by a new product, focusing on the introduction of the minivan. Our model is most closely related to [Gentzkow \(2007\)](#), who develops a demand model that allows for both heterogeneity in consumer preferences and complementarity across products to value online newspapers, and the broader demand estimation literature that allows for complementarity between products ([Hendel, 1999](#); [Dubé, 2004](#); [Thomassen et al., 2017](#); [Allen and Rehbeck, 2019](#); [Compiani, 2022](#); [Iaria and Wang, 2020](#); [Wang, 2024](#); [Ershov et al., 2025](#)). We adapt the framework from [Gentzkow \(2007\)](#) to the rapidly evolving market for subscription AI services and use panel data to separately identify persistent preference heterogeneity and complementarities across providers. We extend the framework by allowing both to evolve over time, and connect the demand estimates to equilibrium price competition. These features allow

us to quantify the surplus generated by paid AI services and study how changing relationships among providers affect competition and the industry’s long-run structure.

2 Data

Our main data is a panel of anonymized credit and debit card transactions from Fiserv, a leading provider of payments and financial services technology solutions. Our sample covers January 2021 through April 2026 and includes information on 114 billion transactions from 719 million cards and 174 million consumers. To limit our sample to instances in which we likely observe individuals’ primary cards, our baseline sample restricts to individual-months with at least 20 transactions. This restriction reduces our sample to 82 million consumers.

The transaction-level data allow us to study within-consumer purchasing patterns over time and across products. This, in turn, enables us to provide novel evidence on competition and substitution among AI products. For each transaction, we observe the purchase date and amount; anonymized card and cardholder identifiers; and the merchant’s name, industry, and location. The data also include consumer demographics, such as age, income group, and ZIP code.

Table 1 reports summary statistics for our sample of consumers. In January 2026, consumers in our sample made an average of 63 transactions and spent an average of \$3,268 over the month. The sample spans the income distribution: 15% of consumers have annual incomes below \$50,000, while 16% have incomes above \$200,000. It also covers all age groups, with 46% of consumers aged 55 or older and 9% aged 34 or younger.

We identify AI-related transactions primarily using merchant names (Appendix A.1 includes additional details). We focus on the four large generative AI providers observed in our data: Anthropic (Claude), Google (Gemini), OpenAI (ChatGPT), and xAI (Grok). Our analysis is restricted to paid consumer accounts, which capture realized consumer demand and provide revealed-preference evidence of consumers’ willingness to pay. They also measure a key source of revenue for providers. Our data do not capture free consumer accounts or enterprise accounts. For some analyses, we also identify transactions for a number of other subscription products, including Grammarly, Chegg, Netflix, Spotify, and the New York Times.

Other Data Sources. We obtain model-capability measures from the Epoch AI Capabilities Index, which assesses performance across a broad range of tasks (Epoch AI, 2025). Figure 1a plots the index for each provider’s frontier model over time. Model capabilities vary substantially over time but relatively little across providers at a given date, showing that AI providers’ frontier models have thus far advanced at similar rates.

We collect product prices from company announcements and archived webpages. We focus on each company’s core entry-level product, which is priced between \$19.99 and \$30 per month for each provider throughout our sample. Figure 1b shows that, across all providers, the vast

majority of those who pay for a product in our data pay less than \$35 per month, consistent with the dominance of these entry-level products.

3 Motivating Evidence: AI Adoption, Multihoming, and Churn

To begin examining the competitive dynamics of the consumer AI industry, we document several new facts about paid AI adoption. We find that multihoming is common but has declined over time, that churn is high relative to other subscription products, and that consumers are particularly likely to move across providers around model releases. These facts motivate and discipline the model developed in Section 4.

3.1 Patterns of AI Adoption

Overall Trends. Consumer use of paid AI products has grown over time, although adoption has varied across providers. Figure 1c shows that by April 2026, 2.9% of consumers paid for at least one of the four major AI products in our sample. Scaling by the number of U.S. credit card holders implies that over 6 million individuals paid for a consumer AI service that month—roughly equal to the combined subscriber bases of the Wall Street Journal and The Economist, or about two-thirds the size of YouTube TV’s subscriber base.

OpenAI accounted for the largest share, with 2.4% of all consumers subscribing. Anthropic, however, experienced particularly rapid growth following the release of Claude 4.6 in February 2026: the share of purchasers tripled from January 2026 to April 2026.

The Geography of Paid AI Adoption. Figure 2 maps paid consumer AI adoption as of April 2026. Adoption is highest along the coasts and in major metropolitan areas: California, Washington, Florida, and the Mid-Atlantic exhibit particularly high rates of AI adoption. By contrast, adoption is much lower in less densely populated areas in the South and Midwest.

Appendix Figures C.1 and C.2 show more localized adoption patterns within Massachusetts and New York City. In Massachusetts, usage is concentrated in the Boston area, reaching as high as 38% in ZIP code 02163 which includes Harvard Business School and its on-campus housing. In New York City, usage is highest in Manhattan—particularly lower and central Manhattan—and northwestern Brooklyn, with a sharp drop-off in the Bronx and in outer Queens and Brooklyn.

Appendix Figure C.3 compares paid AI adoption across different products in the 50 largest commuting zones. Panel (a) shows that OpenAI and Anthropic have very similar spatial patterns, with adoption especially high in San Francisco, Washington, DC, and Austin and lower in Midwestern cities such as Cleveland and St. Louis. Panel (b) shows a similar relationship for xAI, although xAI adoption is relatively less concentrated in the Mid-Atlantic and more concentrated in parts of the South and Southwest, including Austin, Phoenix, and Jacksonville.

Paid AI Adoption by Demographics. Rates of paid AI adoption also differ sharply across demographic groups. Panels (a) and (b) of Figure 3 show that adoption is higher among high-income and younger consumers. As of April 2026, consumers earning over \$200,000 annually are 1.7 times as likely to adopt as those earning less than \$50,000. Similarly, around 3.5% of consumers under 55 pay for AI, compared with 1.9% of those aged 55 and older. Panels (c) and (d) of Figure 3 show that adoption is also strongly related to ZIP-code education and racial composition. In ZIP codes in the highest quartile of college attainment, 4.3% of consumers pay for AI, compared with less than 2% in the lowest quartile. Consumers living in ZIP codes with larger Asian population shares are also substantially more likely to pay for AI.³ Appendix Table C.1 shows broadly similar patterns across products in multivariate regressions.

3.2 Multihoming Rates are High but Decreasing Over Time

To better understand competition in the consumer AI market, we next examine multihoming. Many consumers pay for more than one AI product in a given month, though multihoming has declined over time. These patterns help identify substitutability over time in the model.

Multihoming at OpenAI and Anthropic. To begin, Figure 4a focuses on the two most popular providers as of April 2026, OpenAI and Anthropic. The share of Anthropic users who also pay for OpenAI is high but decreasing: it begins at 44% in June 2024 and falls to 28% in April 2026. The share of OpenAI users who also use Anthropic is smaller, consistent with lower rates of paid Anthropic adoption, but approaches 6% by the end of the sample. The relative stability of these multihoming shares is somewhat at odds with the rapid growth of both OpenAI and Anthropic, as shown in Figure 1c and by the dashed lines in Figure 4a.

Multihoming Index. To more formally assess whether consumers who pay for one AI product are disproportionately likely to also pay for another relative to the unconditional adoption rates, Figure 4b plots the log-odds ratio of joint usage, a commonly used measure of co-adoption. Our measure is:

$$\text{Multihoming Index (OpenAI, Anthropic)} = \log \frac{\text{Share Uses Both} \cdot \text{Share Uses Neither}}{\text{Share Only OpenAI} \cdot \text{Share Only Anthropic}} \quad (1)$$

Intuitively, if adopting OpenAI and Anthropic is statistically independent, then the multihoming index will be zero: both the numerator and denominator will be $p_O p_A (1 - p_O)(1 - p_A)$, where p_O is the probability of using OpenAI and p_A is the probability of using Anthropic. By contrast, positive values imply that the products are used together more often than at random.

³This result is consistent with [Gottfried et al. \(2026b\)](#), a Pew survey finding that nearly 70% of Asian adults use AI tools, compared to roughly half or fewer of Hispanic, Black, or White adults.

Figure 4b shows that the multihoming index (log-odds ratio) remains well above zero for all of the AI pairs throughout the sample, indicating that these products are used together more often than would be expected under independent adoption. However, the index generally declines over time. In June 2024, the multihoming index of OpenAI and Anthropic was 5.2, but this has since fallen to 2.9 in April 2026. This drop is large: the implied odds ratio fell from about 181 to 18, meaning that joint usage went from being about 181 times to 18 times as common, relative to separate usage, as it would be under independent adoption. The figure shows a similar, though less pronounced, decline for other pairs of AI products.

This decline may reflect the fact that AI products have become increasingly general purpose—living up, in a sense, to the promise implied by the product name Chat“GPT”—thereby reducing users’ need to pay for multiple services. For example, if Claude can now serve as a writing assistant in addition to a coding assistant, some consumers may no longer need to subscribe to both Claude and ChatGPT. The extent to which AI products act as substitutes, rather than complements, affects both demand and pricing, as we explore in our model.

3.3 Churn is High and Switching Across Providers is Common

We next look at patterns of churn and substitution across providers. We find that churn is high relative to other subscription markets and that consumers are especially likely to move across providers after model releases. These patterns help identify consumers’ sensitivity to quality and the extent to which their preferences are correlated.

Churn. Panel (a) of Figure 5 shows retention rates for AI products after first use.⁴ OpenAI has the highest retention of all AI products we study, with 77% of people still paying after one month, 65% after three months, and 55% after six months. Anthropic and xAI are lower, but similar to each other: around 65%, 50%, and 40% after one, three, and six months, respectively. Google has the lowest retention rate of the AI products. For benchmarking, the figure also compares these retention rates to the retention rates of non-AI products. Subscriptions like Spotify and Netflix have substantially higher retention, with three-month retention rates above 80%. The New York Times also exceeds all AI providers, with three-month retention after a first purchase of 70%.

Panel (b) of Figure 5 shows how overall retention—the share of all users who, conditional on paying in the last month, paid again this month—changes over time. These rates are generally higher than in panel (a) because they capture the retention of *all* users, not just those first signing up. The figure suggests that overall retention has been relatively stable, although it decreased for xAI in early 2026 as new Anthropic and OpenAI models were released. As of April 2026, 87% of those who paid for OpenAI in one month also did so in the following

⁴To filter out quarterly or annual subscriptions, in this figure, we include only transactions less than \$50.

month. This is somewhat lower at 78% for Anthropic and Google, and 66% for xAI. These rates are again substantially below those of Spotify and Netflix, both of which retain nearly 95% of customers each month in our data.

Model Releases and Switching. We next explore the extent to which consumers switch products around model releases. We again focus on the two largest providers, OpenAI and Anthropic. Figure 6 plots the share of consumers who paid only for Anthropic in the previous month and pay for OpenAI in the current month, a measure of switching or adding an OpenAI product. The share rises sharply around the releases of the GPT-4o and o1 models, with smaller increases around the date ChatGPT added image generation, GPT-5, and Codex Desktop.

Table 2 examines these patterns more systematically using our measure of model-capability from the Epoch Capabilities Index (ECI). We regress the share switching to or adding each provider on indicators for whether that provider released a model that improved its frontier ECI score in the current or previous two months. The sample includes one observation per provider-month for OpenAI and Anthropic, and all specifications include provider fixed effects.

Columns (1)–(3) consider each month separately, while column (4) includes all three. In column (4), a model improvement is associated with a 0.69 percentage point increase in switching or adding that month and a 0.63 percentage point increase the following month—increases of 25% and 23%, respectively, relative to the sample mean. The association is small and statistically insignificant two months after release. Column (5) uses the size of the ECI improvement and also suggests that larger improvements are associated with more switching or adding.

3.4 Exploratory Evidence on the Interaction of AI and Other Spending

We also briefly explore how the release of AI products coincided with changes in demand for other products. Appendix Figure C.4 shows the number of paid users over time around the release of ChatGPT for four potentially related products. Chegg, an online education and homework-help platform, declines substantially following ChatGPT’s release. Midjourney, an AI image-generation service, initially grows before declining sharply after the release of image-generation tools in ChatGPT in October 2023. In contrast to these two, Grammarly, a writing-assistance platform, and Adobe, a provider of creative software, appear to grow or remain relatively stable following ChatGPT’s release. These patterns suggest that while generative AI may substitute for some specific alternative products, others have remained largely unaffected.

To explore consumption patterns more systematically, Appendix Figure C.5 examines changes in spending patterns around each consumer’s first purchase of an OpenAI product. To construct the plot, we match each AI adopter to a non-adopter in the same state-by-age-group-by-income-group cell. AI adopters increase their spending on non-OpenAI software products—which include other AI products—around their first OpenAI purchase. However, their spending in other categories does not appear to strongly differ from that of matched non-adopters.

4 Model

Motivated by our empirical evidence, in this section, we develop a structural model of demand for paid AI services. The model allows us to characterize consumer preferences, their change over time, and implications for market competition. We then use the model to study the equilibrium structure of the consumer AI market and implications of competition on welfare.

4.1 Setup

We model consumers' decisions about which AI services to use as a discrete choice problem in which consumers may choose bundles of services rather than a single service. Our results in Section 3.2 show that consumers frequently multihome across providers. Following Gentzkow (2007), our model distinguishes between two explanations for this multihoming. First, services may be complements, so that combinations of providers are more valuable than the sum of their individual components. Second, multihoming may reflect correlated preference heterogeneity, whereby consumers who strongly prefer one provider also tend to prefer others.

Each consumer i chooses a bundle of AI services from the set of services $\mathcal{J} = 1, \dots, J$. For tractability, we restrict consumers to choosing at most two services. The choice set $\mathcal{C}$ therefore consists of the outside option, the J individual services, and the $\binom{J}{2}$ two-service bundles. We let $c \in \mathcal{C}$ denote a bundle and let $|c|$ denote the number of services contained in bundle c . Under our restriction, $|c| \in 0, 1, 2$.

4.2 Utility

Consumer preferences over AI services are heterogeneous. Consumer i derives utility U_{ict} from choosing bundle $c \in \mathcal{C}$ at time t :

$$U_{ict} = \underbrace{\sum_{j \in c} (\delta_{jt} + \theta_{ij})}_{\text{service-specific utility}} + \underbrace{\Gamma_{ct} + \gamma_{ic}}_{\text{bundle-specific utility}} + \nu_{ict}, \quad (2)$$

where

$$\delta_{jt} \equiv -\alpha p_{jt} + X'_{jt}\beta + \xi_{jt}. \quad (3)$$

The term δ_{jt} captures the mean utility delivered by service j at time t . It consists of a price component, $-\alpha p_{jt}$ with $\alpha > 0$, observed service characteristics X_{jt} , and an unobserved quality component ξ_{jt} . The vector X_{jt} contains observable service characteristics, such as measures of model quality or capability, while ξ_{jt} captures service attributes observed by consumers but not by the econometrician. Consumers differ in their preferences for services through θ_{ij} , which captures consumer i 's time-invariant taste for service j . We allow tastes for services offered by

different providers to be correlated within a consumer, so that multihoming may arise from correlated preference heterogeneity rather than complementarities between services. We assume that these tastes remain fixed over time.

The terms Γ_{ct} and γ_{ic} capture bundle-specific utility that arises when consumers use multiple AI services. We normalize these terms to zero for the outside option and all single-service bundles:

$$\Gamma_{ct} = \gamma_{ic} = 0 \quad \text{for all } c \in \mathcal{C} \text{ such that } |c| \leq 1.$$

Thus, Γ_{ct} and γ_{ic} are only relevant for two-service bundles. The term Γ_{ct} captures the average utility associated with bundle c at time t , while γ_{ic} captures consumer-specific heterogeneity in preferences for that bundle. Allowing Γ_{ct} to vary across bundles and over time permits complementarities between services to evolve as AI technologies and consumer use cases change. Positive values of $\Gamma_{ct} + \gamma_{ic}$ imply that consumer i derives additional utility from consuming the services in bundle c jointly beyond the sum of their standalone utilities, while negative values imply that the services are substitutes.

Finally, ν_{ict} is a consumer-bundle-time-specific utility shock. This term captures transitory factors that affect service choices and rationalizes why consumers do not always choose the same bundle over time. As is standard in the discrete choice literature, we assume that ν_{ict} is independently and identically distributed Type-I Extreme Value across consumers, bundles, and time. Without loss of generality, we normalize the utility of the outside option, corresponding to consuming no AI services, to zero.

4.3 Consumer Choices

Each consumer i chooses the bundle $c \in \mathcal{C}$ that maximizes her utility. The probability that consumer i chooses bundle c , denoted ς_{ict} , is

$$\varsigma_{ict}(\theta_i, \gamma_i) = \frac{\exp\left(\sum_{j \in c} (\delta_{jt} + \theta_{ij}) + \Gamma_{ct} + \gamma_{ic}\right)}{\sum_{c' \in \mathcal{C}} \exp\left(\sum_{j \in c'} (\delta_{jt} + \theta_{ij}) + \Gamma_{c't} + \gamma_{ic'}\right)}. \quad (4)$$

These choice probabilities depend on consumer-specific preferences (θ_i, γ_i) . Aggregating across consumer types, the market share of bundle c is

$$\varsigma_{ct} = \int \frac{\exp\left(\sum_{j \in c} (\delta_{jt} + \theta'_{ij}) + \Gamma_{ct} + \gamma'_{ic}\right)}{\sum_{c' \in \mathcal{C}} \exp\left(\sum_{j \in c'} (\delta_{jt} + \theta'_{ij}) + \Gamma_{c't} + \gamma'_{ic'}\right)} dF(\theta', \gamma'). \quad (5)$$

These bundle shares are the primary objects used in estimation. From the AI-service providers perspective, they care about their total market shares, which depend on their bundle shares. The probability that consumer i purchases service j is obtained by summing the probabilities of

all bundles containing that service,

$$s_{ijt} = \sum_{c \in \mathcal{C}_j} \varsigma_{ict}, \quad (6)$$

where $\mathcal{C}_j$ denotes the set of bundles containing service j . Aggregating across consumers yields the market share of service j ,

$$s_{jt} = \sum_{c \in \mathcal{C}_j} \varsigma_{ct}. \quad (7)$$

Throughout the remainder of the paper, we refer to s_{jt} as product demand and ς_{ct} as bundle demand.

4.4 Supply

The estimated demand system also allows us to study firms' pricing incentives. In our baseline analysis, we assume providers play a static differentiated-product Bertrand pricing game. We then extend the model to account for dynamic learning by doing incentives.

Provider j chooses the monthly price of its service, p_{jt} , taking competitors' prices as given. Assuming constant marginal cost κ_{jt} , provider j solves

$$\max_{p_{jt}} \pi_{jt} = s_{jt}(p_{jt} - \kappa_{jt}). \quad (8)$$

The corresponding first-order condition is

$$s_{jt} + (p_{jt} - \kappa_{jt}) \frac{\partial s_{jt}}{\partial p_{jt}} = 0, \quad (9)$$

or equivalently,

$$p_{jt} - \kappa_{jt} = -\frac{s_{jt}}{\partial s_{jt} / \partial p_{jt}}. \quad (10)$$

Unlike the standard multinomial logit model, a change in the price of service j affects the probability of choosing every bundle containing that service. Consequently, the derivative of product demand is obtained by aggregating the derivatives of the corresponding bundle probabilities:

$$\begin{aligned} \frac{\partial s_{ijt}}{\partial p_{jt}} &= \sum_{c \in \mathcal{C}_j} \frac{\partial \varsigma_{ict}}{\partial p_{jt}} \\ &= -\alpha \sum_{c \in \mathcal{C}_j} \varsigma_{ict} (1 - s_{ijt}) \\ &= -\alpha s_{ijt} (1 - s_{ijt}). \end{aligned} \quad (11)$$

The second equality follows because a change in the price of service j affects the utility of every bundle containing that service, including both the singleton bundle $\{j\}$ and all multihoming bundles. Interestingly, although consumers choose bundles rather than individual services,

aggregating bundle demand to the product level yields the same derivative as in the standard multinomial logit model. Substituting this derivative into the first-order condition yields:

$$\begin{aligned} p_{jt} - \kappa_{jt} &= \frac{s_{jt}}{\alpha \int s_{ijt}(1 - s_{ijt}) dF(\theta, \gamma)} \\ &= - \frac{s_{jt} p_{jt}}{\int s_{ijt} \varepsilon_{ijt} dF(\theta, \gamma)} \end{aligned} \quad (12)$$

where ε_{ijt} denotes consumer i 's own-price elasticity of demand for service j . Thus, despite consumers choosing bundles rather than individual services, the equilibrium markup takes the same form as in the standard Bertrand pricing model. The key difference is that the relevant demand elasticity is computed from product demand, which aggregates substitution across all bundles containing service j . Complementarities enter through the bundle demand system: bundle-specific utility affects the bundle choice probabilities, which in turn determine product demand and its price elasticity.

An important implication of the model is that the degree of complementarity between services directly influences firms' pricing decisions. For example, suppose services j and k become less complementary (or, equivalently, more substitutable) over time, so that Γ_{jkt} decreases. Holding all else constant, consumers become less likely to purchase the bundle $\{j, k\}$, reducing demand for both services. As a result, provider j finds it optimal to lower its price. More generally, changes in complementarities affect equilibrium prices through their impact on product demand and its price elasticity.

4.5 Extension: Dynamic Model Quality

In our baseline version of the model, firms set prices to compete for consumers, taking product quality as fixed. However, we also consider an extension in which product quality evolves over time and depends, in part, on past adoption. Specifically, we introduce learning by doing, whereby providers improve their products as they accumulate users. This mechanism differs slightly from the traditional learning-by-doing framework, in which accumulated experience lowers the marginal cost of production. In our setting, learning by doing instead improves product quality. This extension allows us to study how competition affects not only current prices and adoption, but also the evolution of future product quality.

Consumer Demand. Let ω_{jt} denote the quality of provider j 's model at time t . Model quality is one of the observed characteristics contained in X_{jt} and therefore directly affects consumer utility through equation (3). Let $\beta_\omega > 0$ denote the marginal utility of model quality. Thus, all else equal, improvements in model quality increase current demand. Because current demand determines the number of users from whom a provider learns, higher quality today can increase

adoption, which in turn improves future product quality and future demand.

Evolution of Product Quality. We assume that model quality evolves according to

$$\omega_{jt} = \rho_0 + \rho\omega_{j,t-1} + \lambda s_{j,t-1} + \phi_t + \eta_{jt}, \quad (13)$$

where $s_{j,t-1}$ denotes the share of consumers using provider j in period $t - 1$. The parameter ρ captures persistence in product quality, while λ governs the extent of learning by doing. When $\lambda > 0$, a larger user base today improves the provider's product quality in the following period. The term ϕ_t captures changes in the state of AI technology that are common across providers, while η_{jt} is an idiosyncratic model-quality productivity shock.

Dynamic Pricing. Because current demand affects future product quality, the link between adoption and learning makes the provider's pricing problem dynamic. Given current model qualities ω_t , providers choose prices to maximize the present discounted value of profits. Let $V_j(\omega_t)$ denote provider j 's value function. Provider j solves

$$V_j(\omega_t) = \max_{p_{jt}} \{s_{jt}(p_{jt} - \kappa_{jt}) + \varrho \mathbb{E}_t [V_j(\omega_{t+1})]\}, \quad (14)$$

where $\varrho \in (0, 1)$ denotes the provider's discount factor.

In this framework, providers recognize that their pricing decisions affect the allocation of consumers across providers and, through learning by doing, the future quality of all products in the market. In particular, a lower price for provider j increases its own current demand and therefore its future product quality, while also shifting demand away from competing providers and reducing the experience from which those providers learn. Providers internalize both effects when setting prices. The first-order condition for provider j in the dynamic formulation is:

$$p_{jt} - \kappa_{jt} = \frac{s_{jt}}{\alpha \int s_{ijt}(1 - s_{ijt}) dF(\theta, \gamma)} - \varrho \lambda \sum_{k=1}^J \mathbb{E}_t \left[\frac{\partial V_j(\omega_{t+1})}{\partial \omega_{k,t+1}} \right] \frac{\frac{\partial s_{kt}}{\partial p_{jt}}}{\frac{\partial s_{jt}}{\partial p_{jt}}}. \quad (15)$$

The first term on the right-hand side corresponds to the standard static markup in equation (12). The second term captures the dynamic value of changing the allocation of users across providers through its effect on future model quality. A change in provider j 's price affects not only its own user base and future quality, but also the user bases and future qualities of competing providers. The magnitude of this dynamic incentive depends on the strength of learning by doing, λ , and on how changes in each provider's future quality affect provider j 's continuation value. The own-quality component of the dynamic term gives providers an incentive to lower current prices in order to attract users and improve their future products. At the

same time, a lower price can reduce competitors' future quality by shifting users away from them, providing an additional strategic incentive to compete for users today. When there is no learning by doing ($\lambda = 0$), these dynamic incentives disappear and the pricing condition collapses to the static pricing condition in equation (12).

5 Estimation

We estimate the demand model in two steps. First, we recover consumer preferences via maximum likelihood using panel data on consumer adoption patterns. The estimation identifies the joint distribution of consumer preferences over product characteristics. The panel structure—using repeated observations for a given consumer over time—is valuable for identifying persistent consumer heterogeneity, while substitution patterns over time identify the correlation in unobserved tastes across products and substitutability of products. Second, we relate the estimated preference parameters to observable product characteristics, such as price and model quality, to quantify the sources of consumer heterogeneity.

We then calibrate the supply side of the model. Combining the demand-side estimates with the providers' first-order conditions allows us to recover each provider's marginal cost parameters. Later, we will conduct counterfactuals with our demand-side estimates to study implications of competition on prices and welfare.

5.1 Distribution of Heterogeneous Preferences

Consumers' preferences in equation (2) include a common time- and service-varying component, δ_{jt} ; a heterogeneous individual- and service-specific component, θ_{ij} ; a common time- and bundle-varying component, Γ_{ct} ; and a heterogeneous individual- and bundle-specific component, γ_{ic} . Here, we describe how we parameterize the heterogeneity in consumers' preferences, θ_{ij} and γ_{ic} .

We allow consumers' preferences and the degree of complementarity and substitutability across AI services to be heterogeneous and correlated. We assume that consumers' persistent preference heterogeneity, (θ_i, γ_i) , follows a multivariate normal distribution:

$$\theta_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, \Sigma_\theta), \quad \gamma_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{0}, \Sigma_\gamma), \quad (16)$$

The vector θ_i captures heterogeneous preferences for individual AI services, while γ_i captures heterogeneous preferences for service bundles. We allow preferences for individual services to be correlated within a consumer, so the covariance matrix Σ_θ is unrestricted. For tractability, we assume that preferences for service bundles are independent across bundles and independent

of service-specific preferences:

$$\Sigma_\gamma = \text{diag}(\sigma_{\gamma,c}^2)_{c \in \mathcal{C}_2}, \quad \text{Cov}(\theta_i, \gamma_i) = \mathbf{0}. \quad (17)$$

Allowing for both persistent preference heterogeneity and bundle-specific utility flows enables the model to distinguish between two explanations for the multihoming observed in the data. Consumers may adopt multiple AI services because the services are complements—for example, because ChatGPT increases the value of Claude—or because consumers with strong preferences for one service also tend to have strong preferences for the other. The first mechanism is captured by the bundle-specific preference terms, while the second is captured by correlation in consumers’ service-specific preferences.

5.2 Demand Estimation

Maximum Likelihood Estimation. We estimate the model by maximum likelihood using our quarterly panel of individual bundle choices over the period 2024Q1-2026Q1. Let d_{ict} equal one if consumer i chooses bundle $c \in \mathcal{C}$ at time t and zero otherwise, and let $d_i = \{d_{ict}\}_{c \in \mathcal{C}, t=1, \dots, T}$ denote consumer i ’s observed sequence of choices. The parameters to be estimated are $\Theta = (\{\delta_{jt}\}, \{\Gamma_{ct}\}, \Sigma)$, where Σ denotes the parameters corresponding to the distribution of preference heterogeneity $F(\theta, \gamma)$. We treat the mean utilities δ_{jt} and mean bundle utilities Γ_{ct} as parameters to be estimated directly, and in a second stage decompose δ_{jt} into its price, characteristic, and unobserved quality components according to equation (3).

Conditional on a consumer’s persistent tastes (θ_i, γ_i) , the utility shocks ν_{ict} are independently and identically distributed over time. Consumer i ’s choices are therefore independent across periods conditional on (θ_i, γ_i) , and the probability of observing the choice sequence d_i for a consumer with tastes (θ, γ) is the product of the period-by-period choice probabilities,

$$\Pr(d_i \mid \theta, \gamma; \Theta) = \prod_{t=1}^T \prod_{c \in \mathcal{C}} s_{ict}(\theta, \gamma)^{d_{ict}}. \quad (18)$$

Because the econometrician does not observe (θ_i, γ_i) , the likelihood contribution of consumer i integrates over the taste distribution,

$$L_i(\Theta) = \int \prod_{t=1}^T \prod_{c \in \mathcal{C}} s_{ict}(\theta, \gamma)^{d_{ict}} dF(\theta, \gamma). \quad (19)$$

We describe identification further below (Section 5.2.2); however, the structure of equation (19) highlights the role of the panel in identifying persistent heterogeneity. The integral over $F(\theta, \gamma)$ enters outside the product over time, so persistent tastes induce correlation in a consumer’s choices across periods, while the transitory shocks ν_{ict} do not. Intuitively, a con-

sumer who repeatedly chooses the same service or the same bundle reveals a large persistent taste component, whereas switching across bundles over time is attributed to the transitory shock. It is this within-consumer persistence in choices that separately identifies the distribution of (θ_i, γ_i) from the variance of ν_{ict} , and, together with observed bundling patterns, distinguishes correlated preference heterogeneity from complementarities Γ_{ct} .

The integral in equation (19) does not admit a closed form, so we approximate it by simulation. We draw R realizations (θ^r, γ^r) , $r = 1, \dots, R$, from $F(\theta, \gamma)$ and form the simulated likelihood contribution

$$\hat{L}_i(\Theta) = \frac{1}{R} \sum_{r=1}^R \prod_{t=1}^T \prod_{c \in \mathcal{C}} s_{ict}(\theta^r, \gamma^r)^{d_{ict}}. \quad (20)$$

We estimate Θ by maximizing the simulated log-likelihood $\sum_i \log \hat{L}_i(\Theta)$.

Implementation Details. We estimate the model on individuals who have at least 20 transactions in each of the nine quarters from 2024Q1-2026Q1 ($T = 1, \dots, 8$ from 2024Q2-2026Q1 taking initial shares in 2024Q1 as given). A consumer is classified as having adopted the AI bundle in quarter t if they used all the products in that bundle at any point in the quarter. We focus on the four largest AI providers in our data (OpenAI, Anthropic, xAI, and Google) and allow multi-homing up to two products, so there are $c = 10$ possible bundles that are non-empty.⁵

As a result, there are 96 parameters which we estimate iteratively using L-BFGS with increasingly lower tolerance. In particular, we estimate average preference heterogeneity δ_{jt} (4×8), the lower-diagonal of its variance-covariance matrix Σ_θ (10), complementarities from using multiple AI services Γ_{ct} (6×8), and its diagonal variance matrix Σ_γ (6). Appendix B.1 contains additional details on implementation.

5.2.1 Recovering Price and Quality Sensitivity

Maximum likelihood estimation allows us to recover the mean service and bundle utilities as well as the distribution of preference heterogeneity. We further decompose the mean product-specific utility into its price, characteristics and unobserved quality components by regressing the recovered utilities δ_{jt} on observable characteristics

$$\hat{\delta}_{jt} = -\hat{\alpha} p_{jt} + X'_{jt} \hat{\beta} + \hat{\xi}_{jt}. \quad (21)$$

We estimate α and β in Equation 21 using linear regression. Since OpenAI, Anthropic, and Google held monthly subscription prices constant throughout our sample, the only within-provider pricing variation comes from xAI. To isolate this variation, we instrument for price using a dummy equal to one for xAI observations from Q1 2025 onward — the period following

⁵In our sample, fewer than 0.05% of consumers have ever used more than two products.

xAI’s price increase from \$16 to \$30 per month.⁶

5.2.2 Identification

The parameters are identified from consumers’ observed choice behavior. While all parameters are jointly identified in the model, it is useful to provide intuition for the patterns in the data that are most informative about each set of parameters.

The mean utility terms, δ_{jt} and Γ_{ct} , are identified primarily from average choice frequencies. For example, if ChatGPT has a large market share, this implies a high value of $\delta_{\text{ChatGPT},t}$. Similarly, if the ChatGPT–Claude bundle is chosen more frequently than would be predicted based on the individual market shares of ChatGPT and Claude, this suggests a higher value of Γ_{ct} , all else equal.

The distribution of preference heterogeneity, $(\Sigma_\theta, \Sigma_\gamma)$, is identified from persistent within-consumer choice behavior. For example, if a consumer repeatedly chooses ChatGPT over time, this persistence cannot be fully explained by δ_{jt} because δ_{jt} is common across consumers. Instead, such behavior is indicative of persistent heterogeneity in preferences. More generally, the diagonal elements of Σ_θ and Σ_γ are identified from the extent to which consumers repeatedly make similar choices over time.

The off-diagonal elements of Σ_θ are identified from co-adoption and substitution patterns across AI services. For example, if consumers who currently use or previously used Claude are systematically more likely to adopt ChatGPT, this implies a positive covariance between preferences for the two services, i.e.,

$$\text{Cov}(\theta_{i,\text{ChatGPT}}, \theta_{i,\text{Claude}}) > 0. \quad (22)$$

Empirically, a central challenge is separating the interaction terms, Γ_{ct} , from correlated preference heterogeneity, $(\Sigma_\theta, \Sigma_\gamma)$. For example, frequent selection of the ChatGPT–Claude bundle could indicate that the products are complements, but could also reflect persistent unobserved heterogeneity: some consumers may simply have high demand for AI products. We distinguish these channels using panel variation, following the logic of [Gentzkow \(2007\)](#). The key identifying assumption is that the correlated component of unobserved preferences is persistent over time. Under this assumption, if two AI products are complements, a consumer should be more likely to use both products in the same periods and neither in others. By contrast, if multi-homing is driven by correlated taste heterogeneity, high-AI-type consumers will have persistently high demand for AI tools, but may switch between single-homing and multi-homing as relative product quality changes, including around model releases, as suggested by Figure 6.

⁶In December 2024, xAI raised prices of Premium+ from \$16 to \$22. In February 2025, after the release of Grok 3, they introduced SuperGrok at \$30 and raised the price of Premium+ to \$40 (see, e.g., [Mehta, 2025](#)).

5.3 Supply Calibration

Using the estimated parameters, observed market shares and prices, and the providers' first-order conditions (eq. 12), we calibrate the supply side of the model, which consists of the providers' marginal costs of providing AI services, κ_{jt} . Rearranging eq. (12), we solve for marginal costs as

$$\kappa_{jt} = p_{jt} - \frac{s_{jt}}{\hat{\alpha} \int s_{ijt}(1 - s_{ijt}) d\hat{F}(\theta, \gamma)}. \quad (23)$$

We directly observe prices and market shares in the data. The remaining terms depend on the model estimates: we compute the individual choice probabilities s_{ijt} at the estimated parameters and approximate the integral by simulation, drawing from $\hat{F}(\theta, \gamma)$.

6 Results

Our model results suggest that paid consumer demand for AI is sensitive to quality and prices, that our multihoming patterns in Section 3 are driven by correlated preferences rather than complementarities in the products themselves, and that AI products are becoming more substitutable over time. Substantial preference heterogeneity implies that the welfare effects of AI are concentrated among heavy AI users.

6.1 Consumer Preferences

We report our estimates of consumer preferences in Table 3. Throughout, we convert utility estimates into willingness-to-pay by dividing by the estimated price coefficient $\hat{\alpha}$. Panel (a) displays our mean preference parameters. Consistent with the overall market shares, the results suggest that the typical consumer prefers OpenAI products relative to other alternatives, a preference that has remained stable over time. For example, our estimates imply that the average consumer was willing to pay a \$10 premium per month for OpenAI products relative to Anthropic products at the end of the sample in 2026Q1.

The bottom row of panel (a) displays the standard deviations of preference heterogeneity. The results indicate that there is substantial heterogeneity in how individuals value AI, which is unsurprising given the variation in adoption patterns. The standard deviations range from 5.5 to 6 across providers, which is substantial relative to the mean valuations (on average 34% of the size). Consumers in the top quartile of the preference distribution for ChatGPT are willing to pay \$37 more per month for ChatGPT than consumers in the bottom quartile.

In addition to heterogeneity, our estimates suggest that consumers are relatively sensitive to prices and model quality. The first row of panel (b) displays our estimates of consumers' price sensitivity and how they value AI capabilities. In each specification, we estimate a negative and significant monthly price coefficient. Our instrumental variables estimate is -0.22 which

implies an average price elasticity of -4.6; that is a 10% increase in price lowers demand by 46%. The second row of panel (b) displays the quality coefficient, which suggests that moving from ChatGPT’s Model 4 in March 2023 to Model 5.5 in April 2026 increased the value to consumers by \$25 per month. In other words, consumers are indifferent between GPT-4 and paying a \$25 per month premium to access Model GPT-5.5.

Panel (c) displays the correlations in preference heterogeneity. The results suggest that a consumer’s preferences for AI are positively correlated within an individual. The correlation in preferences between ChatGPT and Claude is 0.89. Customers who value xAI and Google tend to have lower preferences for the other AI providers, consistent with both serving a distinct consumer segment from general-purpose AI assistants. (Gemini, for example, is often bundled with GoogleOne products). As a result, the high levels of multihoming seen in Figure 4 are likely due to correlated preferences rather than complementarities.

Indeed, panel (d) displays the mean and standard deviation of the complementarity between different AI providers over time. The results suggest that products like OpenAI and Google were complements at the beginning of the sample, but have since become substitutes. OpenAI and Anthropic were always substitutes, even more so towards the end of the sample. Holding preferences fixed at their values at the start of the sample, so that the comparison isolates the complementarity channel from correlated tastes, the probability that users with Claude would also have ChatGPT decreased by 83% over the sample. This decline in complementarity increases competition: it reduces demand for each service and, all else equal, puts downward pressure on equilibrium prices, as will be illustrated in Section 7.

6.2 Welfare

Our estimated model implies that the welfare effects of AI are heterogeneous across consumers, and have increased over time as models have improved. Consistent with the substantial preference heterogeneity we estimate in Table 3 panel (a), the surplus from AI is highly concentrated: consumers in the top decile of the OpenAI taste distribution capture over 97% of consumer surplus. As AI models improved over the two years in our sample, total consumer surplus has increased over five-fold.

Estimation. The consumer surplus from adding an AI provider (e.g., OpenAI) for consumer i in period t is $\Delta CS_{it}^O = W_{it}(\mathcal{C}_t) - W_{it}(\mathcal{C}_t^{-O})$ where $W_{it}(\mathcal{C}_t)$ is the expected indirect utility from access to choice set $\mathcal{C}_t$ (in dollar terms). $\mathcal{C}_t$ is the full choice set and $\mathcal{C}_t^{-O} = \{c \in \mathcal{C}_t : \text{OpenAI} \notin c\}$ removes all choices with OpenAI. Since the utility shocks, ν_{ict} , in equation 2 are distributed Type-I Extreme Value, the expected indirect utility from access to choice set $\mathcal{C}_t$ in dollars is:

$$W_{it}(\mathcal{C}_t) = \frac{1}{\alpha} \left(\log \sum_{c \in \mathcal{C}_t} \exp(u_{ict}) + K \right). \quad (24)$$

where $u_{ict} = \sum_{j \in c} (\delta_{jt} + \theta_{ij}) + \Gamma_{ct} + \gamma_{ic}$ is the deterministic component of consumer i 's utility from choosing bundle c in period t and K is the Euler-Mascheroni constant. Therefore the change in consumer surplus is:

$$\Delta CS_{it}^O = \frac{1}{\alpha} \left[\log \left(\sum_{c \in \mathcal{C}_t} \exp(u_{ict}) \right) - \log \left(\sum_{c \in \mathcal{C}_t^O} \exp(u_{ict}) \right) \right]. \quad (25)$$

Results. Figure 7 plots total consumer welfare from paid AI products. The welfare gain from AI has risen five-fold in the two years prior to 2026Q1, from \$0.20 per quarter to over \$1.09 per quarter. Since OpenAI has the largest market share among paid consumers, 74% of the welfare gain from consumer AI products can be attributed to OpenAI. Recent growth in Anthropic implies that by 2026Q1, welfare from Anthropic has risen to \$0.13 per quarter.

The average increase in welfare masks substantial heterogeneity because most consumers do not currently subscribe to an AI service. Dividing each average welfare gain by the share of consumers paying for any AI service, the normalized quarterly gains are \$35 for all AI products in 2026Q1, with \$26 attributable to OpenAI and \$4 attributable to Anthropic.

6.3 Supply Parameters

Table 4 shows our estimates of markups, estimated with and without dynamic model quality. With a static supply-side (Equation 23), the model implies that OpenAI charges a markup of approximately \$12 per month (Lerner index of 59%), whereas Anthropic charges a markup of \$9 per month. Anthropic charges a lower markup, in part, because OpenAI and Anthropic are substitutes.

Our model with dynamic quality incentives suggests an even more competitive market than in the baseline model. Full details on our dynamic calibration are in Section B.2. With dynamic model quality, Column 2 shows that implied markups are lower for all providers, but significantly so for OpenAI, falling from an implied markup of \$12 per month to \$4 per month. This is because OpenAI has the largest customer base, so if the dynamic quality incentive is meaningful, it would be strongest for OpenAI. Overall, our dynamic model estimates suggest that marginal costs account for roughly 70-80% of revenues. Our estimates are an upper bound on profitability as they ignore fixed operating costs and research and development expenditures.

7 Counterfactuals and Implications for Competition

We use our model to examine how product substitutability and preference heterogeneity shape competition. We find that equilibrium prices would be substantially higher if AI products were, on average, complements rather than substitutes or if consumers' preferences across products

were independent rather than positively correlated. Because Anthropic’s and OpenAI’s products are particularly close substitutes and consumers’ preferences for them are highly correlated, Anthropic plays an important role in constraining market power despite its relatively small current market share. A counterfactual without Anthropic indicates that a meaningful share of the resulting welfare loss comes from higher prices charged by the remaining providers (in addition to the direct loss of access to Anthropic’s products).

7.1 Effects of Substitutability on Competition

The increased substitutability between AI products has several important implications for market structure. As in [Gentzkow \(2007\)](#), even though multihoming rates are unconditionally high, our model suggests that this is largely from correlated preference heterogeneity rather than AI products exhibiting strong complementarities. To illustrate the impact of substitutability on prices and market shares, we consider a counterfactual in 2026Q1 where $\Gamma_{counter} = 0$ and $\Gamma_{counter} = |\Gamma|$. Without substitutability, consumers respond by being more likely to bundle AI products together. This increases the share of consumers using AI, but providers dampen this effect by raising prices since there is more market power.

Table 5 shows that if $\Gamma = 0$, prices would be higher for all providers, and in net, shares of most providers would be higher than the baseline (Figure 8 plots counterfactual shares over time). Anthropic is a particularly interesting case. In this counterfactual, Anthropic raises their prices from \$20 per month to \$26 per month. However, even with this price change, Anthropic’s market share nearly doubles, from 0.50% to 0.96%. OpenAI, GoogleAI and xAI experience smaller changes in prices and shares. The disproportionate increase in Anthropic’s market shares and prices are likely because OpenAI (which has the largest baseline market share) is by far the most substitutable with Anthropic: $\Gamma_{OpenAI, Anthropic} = -6.44$, compared with $\Gamma_{OpenAI, Google} = -0.90$ and $\Gamma_{OpenAI, xAI} = -0.04$.

Our results are stronger if AI products were all complements rather than neither complements nor substitutes. Relative to the baseline, we find that with $|\Gamma|$, providers increase their prices by \$6-\$12 per month. The share of consumers paying for OpenAI and Anthropic rises to 3.02% and 2.33%, respectively, representing a percentage increase of 16% and 364%.

Overall, these results suggest that as AI providers have evolved into general-purpose technologies, their products have become increasingly substitutable. As a result, the marketplace has become more competitive, putting downward pressure on prices.

7.2 Impact of Correlated Preferences on Competition

Product differentiation is an important source of market power. Although we find substantial heterogeneity in consumers’ overall preferences for AI, there is limited differentiation in their preferences across providers. As shown in panel (c) of Table 3, preferences for AI products are

highly correlated. Conditional on how much a consumer likes AI overall, the consumer tends to have similarly favorable preferences for all AI providers.

To illustrate the importance of correlated preferences for competition, Table 5 considers a counterfactual in which preferences are uncorrelated, so the off-diagonal entries of Σ_θ are zero. Under this counterfactual, prices would be unambiguously higher in 2026Q1. In particular, all providers modestly increase their prices, with changes ranging from 2.1% for OpenAI, which has the smallest price response, to 8.7% for xAI, which has the largest price response.

The directional change in market shares differs across providers. The market shares of Anthropic and xAI increase by 28.8% and 29.8%, respectively, whereas the market shares of OpenAI and GoogleAI decrease slightly—by less than 3%. Intuitively, Anthropic’s preferences are the most strongly correlated with those of the other three providers. Its market share therefore increases substantially when consumers who prefer Anthropic no longer have strong preferences for its competitors. xAI also experiences a sizable increase in market share because preferences for xAI and GoogleAI are highly correlated in the baseline and GoogleAI is cheaper. As a result, GoogleAI exerts substantial competitive pressure on xAI in the baseline, which is removed in the counterfactual.

7.3 Welfare Effects of Entry

Our estimates imply that AI products are close substitutes, creating competitive pressure among providers. A central implication is that the presence of an individual provider—even one that is not the market leader—may meaningfully constrain the prices charged by its rivals. To quantify this competitive discipline, we consider a counterfactual in which Anthropic is removed from the market in 2026Q1 and allow the remaining providers to reoptimize their prices. Anthropic is a particularly informative case because it is a close substitute for OpenAI, while its users also have overlapping preferences for xAI and GoogleAI. The counterfactual therefore illustrates how the loss of one provider propagates through equilibrium prices, market shares, and consumer welfare.

Panel (a) of Figure 9 shows the effect of Anthropic’s counterfactual exit on equilibrium prices. OpenAI and xAI increase their prices by 2.0% and 1.1%, respectively, whereas GoogleAI changes its price by -4.4%. Thus, removing Anthropic does not uniformly increase the market power of every remaining provider. Instead, its effects depend on the patterns of substitution and preference correlation across products. Although the price responses are modest in magnitude, they are notable given Anthropic’s small market share. As we show below, these price changes also affect paid adoption and contribute meaningfully to the change in consumer surplus.

Panel (b) shows that the changes in equilibrium market shares are similarly heterogeneous. Interestingly, OpenAI’s market share changes by -0.7%, even though its closest substitute has been removed from the market. This seemingly counterintuitive result reflects OpenAI’s en-

ogenous pricing response. Anthropic’s removal relaxes the competitive constraint on OpenAI, allowing it to raise its price. This price increase more than offsets the direct inflow of consumers who would otherwise switch from Anthropic to OpenAI, causing OpenAI’s equilibrium market share to fall.

GoogleAI’s market share also falls, by -7.0%, whereas xAI’s market share increases by 17.1%. xAI’s disproportionate response reflects the distribution of consumers’ preferences for bundling. Although OpenAI and Anthropic are the most substitutable pair on average, as captured by Γ , the individual-level estimates of Γ for xAI and Anthropic have very low dispersion. Consequently, most Anthropic users view xAI as a net substitute, and removing Anthropic shifts many of these users toward xAI despite xAI’s higher equilibrium price. GoogleAI and Anthropic, by contrast, are not strong net substitutes. Moreover, xAI’s expansion intensifies competition with GoogleAI because preferences for these two products are highly correlated, reducing GoogleAI’s market power and inducing it to lower its price.

Accounting for these supply-side responses is important for evaluating welfare. The direct loss of access to Anthropic understates the total welfare effect because the remaining providers endogenously adjust their prices. In particular, the price increases by OpenAI and xAI amplify the direct welfare loss from removing Anthropic. Panel (c) shows that, when expressed per AI user, removing Anthropic reduces welfare by \$5.05 per quarter. This general-equilibrium welfare loss is 21% larger than the partial-equilibrium loss reported in Section 6.2.

Overall, the counterfactual demonstrates that competition in the AI market depends not only on the number and market shares of providers, but also on the structure of substitution among their products. Because consumer preferences are highly correlated across providers and AI products are strong substitutes, Anthropic impacts its rivals’ prices even though consumers frequently use multiple products. Removing Anthropic therefore generates welfare losses beyond the direct loss of access to its product, as the pricing responses of the remaining providers amplify the initial harm.

8 Conclusion

We document new patterns in consumer demand for paid generative AI services. Despite AI’s prominence in public and business discourse, paid consumer adoption remains relatively modest, although it grew rapidly toward the end of our sample. In April 2026, 2.4% of consumers paid for an OpenAI product, making OpenAI the most widely adopted provider. Paid adoption of Anthropic products also increased substantially during the final months of our sample. Adoption is concentrated among younger and higher-income consumers. We also find that paid demand exhibits limited persistence and that multihoming across providers is common, particularly early in the sample, although multihoming has declined over time.

Our structural model provides additional insight into these patterns. We estimate substan-

tial heterogeneity in consumer preferences, but that preferences for general-purpose AI services are highly correlated across providers. At the same time, the relationship between OpenAI and Anthropic has evolved: the two services have become increasingly substitutable over time. This shift is consistent with the growing overlap in their capabilities.

Our supply-side model shows how these demand patterns shape competition. Increasing substitutability and overlap among services intensify price competition. In particular, consumers view services from Anthropic and OpenAI as close substitutes. These competitive forces affect both the total welfare generated by paid AI services and the division of surplus between consumers and providers. The resulting limits on firms' ability to sustain market power also bear directly on firm valuations throughout the industry.

References

- ALLEN, R. and REHBECK, J. (2019). Identification with additively separable heterogeneity. *Econometrica*, **87** (3), 1021–1054.
- ATHEY, S. and MORTON, F. (2025). *Artificial Intelligence, Competition, and Welfare*. NBER Working Paper 34444.
- BABINA, T., HE, A. X. and JIANG, R. (2026). *Canaries in the gold mine: Early productivity gains from artificial intelligence creating organization capital*. NBER Working Paper 35684.
- BANK OF AMERICA INSTITUTE (2026). *Not Quite mAIstream: A Consumer AI Adoption Profile*. Tech. rep., Bank of America Institute, march 30, 2026.
- BERRY, S., LEVINSOHN, J. and PAKES, A. (1995). Automobile prices in market equilibrium. *Econometrica*, **63** (4), 841–890.
- , — and — (2004). Differentiated products demand systems from a combination of micro and macro data: The new car market. *Journal of Political Economy*, **112** (1), 68–105.
- BICK, A., BLANDIN, A. and DEMING, D. J. (2026a). The rapid adoption of generative AI. *Management Science*.
- , —, —, FUCHS-SCHÜNDELN, N. and JESSEN, J. (2026b). Mind the gap: AI adoption in Europe and the US. *Working Paper*.
- , —, — and SCHUMACHER, T. R. (2026c). *What Work Does Generative AI Do?* NBER Working Paper 35677.
- BOARD OF GOVERNORS OF THE FEDERAL RESERVE SYSTEM (2026). *Economic Well-Being of U.S. Households in 2025*. Tech. rep., Board of Governors of the Federal Reserve System.
- BONNEY, K., BREAUX, C., DINLERSOZ, E., FOSTER, L., HALTIWANGER, J. and PANDE, A. (2026). *The Microstructure of AI Diffusion: Evidence from Firms, Business Functions, and Worker Tasks*. CES Working Paper 26-25, U.S. Census Bureau, Center for Economic Studies.
- BRYNJOLFSSON, E., COLLIS, A., EGGERS, F., KAZINNIK, S. and NGUYEN, D. (2026). *What is Generative AI Worth*. Tech. rep., working Paper, Available at SSRN.
- , LI, D. and RAYMOND, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*, **140** (2), 889–942.
- CAOUI, E. H. (2026). Estimating consumer preferences for llms: Evidence from lmarena.
- CHATTERJI, A., CUNNINGHAM, T., DEMING, D., HITZIG, Z., ONG, C., SHAN, C. and WADMAN, K. (2025). *How people use ChatGPT*. NBER Working Paper 34255.


- CHEN, F. and STRATTON, J. (2026). Artificial intelligence in the firm: Bottlenecks in software production. Working Paper.
- CHEN, Z. and CHAN, J. (2024). Large language model in creative work: The role of collaboration modality and user expertise. *Management Science*, **70** (12), 9101–9117.
- CHOUKHMANE, T., DE SILVA, T., LIN, W. and AKUZAWA, M. (2026). *AI Financial Advice: Supply, Demand, and Life Cycle Implications*. NBER Working Paper 35574.
- COMPIANI, G. (2022). Market counterfactuals and the specification of multiproduct demand: A nonparametric approach. *Quantitative Economics*, **13** (2), 545–591.
- CUI, K. Z., DEMIRER, M., JAFFE, S., MUSOLFF, L., PENG, S. and SALZ, T. (2024). *The productivity effects of generative AI: Evidence from a field experiment with GitHub Copilot*. Tech. rep., working Paper.
- DELAPREZ, Y. and HORTAÇSU, A. (2026). The welfare impact of delegating choice to large language models. Working paper, available at SSRN.
- DEMIRER, M., FRADKIN, A., TADELIS, S. and PENG, S. (2025). *The Emerging Market for Intelligence: Pricing, Supply, and Demand for LLMs*. NBER Working Paper 34608.
- DILLON, E. W., JAFFE, S., IMMORLICA, N. and STANTON, C. T. (2025). *Shifting work patterns with generative AI*. NBER Working Paper 33795.
- DUBÉ, J.-P. (2004). Multiple discreteness and product differentiation: Demand for carbonated soft drinks. *Marketing Science*, **23** (1), 66–81.
- EPOCH AI (2025). Epoch Capabilities Index. Accessed: 30 Sep 2026.
- ERSHOV, D., MARCOUX, M., ORR, S. and LALIBERTÉ, J.-W. P. (2025). Estimating complementarity with large choice sets: An application to mergers. *RAND Journal of Economics*, **56** (4), 689–707.
- GAVIN, W. (2026). OpenAI is now bringing in \$2 billion a month and 3 more highlights from its latest update. *MarketWatch*.
- GENTZKOW, M. (2007). Valuing new goods in a model with complementarity: Online newspapers. *American Economic Review*, **97** (3), 713–744.
- GOTTFRIED, J., BISHOP, W., ANDERSON, M., FAVERIO, M., PARK, E. and MCCLAIN, C. (2026a). *Americans and AI 2026: Chatbots, Smart Devices and Views on Impact*. Tech. rep., Pew Research Center.

- , —, —, —, — and — (2026b). Racial and ethnic differences in how adults use and view AI. Pew Research Center, accessed August 31, 2026.
- HAGIU, A. and WRIGHT, J. (2023). Data-enabled learning, network effects, and competitive advantage. *RAND Journal of Economics*, **54** (4), 638–667.
- HAMPOLE, M., PAPANIKOLAOU, D., SCHMIDT, L. D. and SEEGMILLER, B. (2025). Artificial intelligence and the labor market.
- HANDA, K., TAMKIN, A., MCCAIN, M. *et al.* (2025). Which economic tasks are performed with AI? evidence from millions of claude conversations.
- HENDEL, I. (1999). Estimating multiple-discrete choice models: An application to computerization returns. *Review of Economic Studies*, **66** (2), 423–446.
- IARIA, A. and WANG, A. (2020). *Identification and Estimation of Demand for Bundles*. CEPR Discussion Paper 14363.
- KEMP, A. (2025). AI use at work rises.
- LIN, B. (2025). Anthropic lands Cognizant as enterprise AI customer. *The Wall Street Journal*.
- LU, F., LI, S. and SHI, Y. (2025). Robo-advising meets large language models: Educating investors on alpha and beta of mutual funds and stocks. Working paper, available at SSRN.
- MEHTA, I. (2025). X doubles its premium+ plan prices after xAI releases grok 3. *TechCrunch*.
- NEVO, A. (2000). A practitioner’s guide to estimation of random-coefficients logit models of demand. *Journal of Economics & Management Strategy*, **9** (4), 513–548.
- NOY, S. and ZHANG, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, **381** (6654), 187–192.
- OTIS, N., CLARKE, R., DELECOURT, S., HOLTZ, D. and KONING, R. (2024). The uneven impact of generative AI on entrepreneurial performance. Working Paper, Available at SSRN.
- PARHAM, A. (2026). OpenAI enterprise revenue tops consumer for first time: \$40 billion arr two quarters early. *Tech Times*.
- PETRIN, A. (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy*, **110** (4), 705–729.
- PNC ECONOMICS RESEARCH (2026). *Consumer Health Check: Spending on Cruise Control*. Tech. rep., PNC.

- PRUFER, J. and SCHOTTMÜLLER, C. (2021). Competing with big data. *Journal of Industrial Economics*, **69** (4), 967–1008.
- SARKAR, S. K. (2025). AI agents, productivity, and higher-order thinking: Early evidence from software development. Working paper, available at SSRN.
- SONG, F., AGARWAL, A. and WEN, W. (2026). The impact of generative AI on collaborative open-source software development: Evidence from github copilot. *Information Systems Research*.
- THOMASSEN, Ø., SMITH, H., SEILER, S. and SCHIRALDI, P. (2017). Multi-category competition and market power: A model of supermarket pricing. *American Economic Review*, **107** (8), 2308–2351.
- WANG, A. (2024). A BLP demand model of product-level market shares with complementarity, working paper, University of Warwick.
- YOTZOV, I., BARRERO, J. M., BLOOM, N., BUNN, P., DAVIS, S. J., FOSTER, K. M., JALCA, A., MEYER, B. H., MIZEN, P., NAVARRETE, M. A. *et al.* (2026). Firm data on AI. Federal Reserve Bank of Atlanta Working Paper 2026-3.


Figures and Tables

Figure 1: AI Model Capability, Prices, and Paid Adoption




Notes: Panel **a** shows frontier AI model capability over time for Anthropic, Google, OpenAI, and xAI, as measured by the Epoch AI Capabilities Index (ECI). Lines show each provider's best available model score by month; dots mark months in which its frontier score improved. Panel **b** shows the distribution of subscription prices paid across AI providers. Each series reports the share of a provider's subscribers in one of four monthly price bins: below \$35, [\$35,\$100), [\$100,\$200), and \$200 or more. Panel **c** shows the share of consumers paying for AI products in our baseline sample, described in Section 2.

Figure 2: Geography of Paid AI Adoption



Note: Figure 2 maps the share of consumers paying for AI across commuting zones in the United States as of April 2026. Darker colors indicate a higher rate of paid AI-product adoption. Gray commuting zones have fewer than 100 individuals in our baseline sample.


Figure 3: Adoption of Paid AI Products by Demographics




Note: Figure 3 shows the share of consumers paying for AI products over time by demographic group. Panels (a) and (b) report trends by income and age, respectively. Panel (c) groups consumers by the college-educated share of their ZIP code, with Q1 denoting the lowest share and Q4 the highest. Panel (d) constructs race-specific trends by weighting each consumer by the racial composition of their ZIP code.

Figure 4: Multihoming Across AI Products

(a) Multihoming at Anthropic and OpenAI




(b) Multihoming index over time




Note: Panel (a) shows the share of each provider’s paid subscribers who also subscribe to the other provider (solid lines, left axis) and each provider’s paid-subscriber adoption per 100,000 individuals (dashed lines, right axis). Panel (b) plots the multihoming index, defined as the log-odds ratio of joint usage in Equation (1), for each pair of AI products over time. Positive values indicate that products are used together more often than would occur under statistical independence.

Figure 5: Retention of AI Subscriptions

(a) Retention Rates After First Use




(b) Month-to-Month Retention Over Time




Note: Figure 5 shows measures of AI consumer retention. Panel (a) shows retention rates for AI products up to one year after first use, benchmarked against selected non-AI subscription products. Panel (b) shows month-to-month retention over time. This is the share of users who, conditional on subscribing last month, subscribe again this month. To filter out quarterly and annual subscriptions, we restrict to transactions below \$50. In panel (a), we pool all cohorts observed for at least M months to calculate retention at month M . For Google AI in panel (a), we begin the curve at the second observed month of use because we identify subscribers using transactions in consecutive months (see Appendix A.1).

Figure 6: Switching to OpenAI from Anthropic




Note: Figure 6 shows the monthly share of consumers who pay for an OpenAI product among those who paid only for an Anthropic product in the previous month. Vertical lines indicate OpenAI model or feature releases, with releases after the middle of the month assigned to the following month.

Figure 7: Consumer Welfare from AI Over Time




Note: Figure 7 plots estimated consumer welfare from paid AI products over time, expressed in dollars per consumer and averaged across all consumers, including those who do not currently pay for AI. Welfare is computed as the loss in expected indirect utility from removing products from the choice set, as described in equation (25). The green line shows the welfare loss from removing all AI products; the blue and orange lines show the welfare loss from removing OpenAI and Anthropic, respectively.

Figure 8: Counterfactual Market Shares with No Complementarities



Note: Figure 8 compares observed AI product market shares to counterfactual shares in which deterministic bundle-specific utility is set to zero ($\Gamma_{ct} = 0$), holding all other estimated parameters fixed. The counterfactual isolates the effect of complementarity and substitutability on product demand.

Figure 9: Impacts of Removing Anthropic



Note: Figure 9 shows the impacts of removing Anthropic in 2026Q1. Panel (a) shows the price response, panel (b) the change in shares, and panel (c) the welfare loss for both all consumers and scaled by the share using AI.

Table 1: Baseline Sample Summary Statistics

	Jan 2025	Jul 2025	Jan 2026
Total consumers	32,490,671	33,950,133	30,691,583
Total transactions	1,850,075,422	2,132,764,052	1,922,919,178
Avg transactions per consumer	57	63	63
Avg spend per consumer (\$)	2,932	3,256	3,268
Age: <35 (%)	9	8	9
Age: 35-54 (%)	32	32	32
Age: 55-64 (%)	22	22	22
Age: 65+ (%)	25	25	25
Income: <\$50k (%)	15	15	15
Income: \$50k-100k (%)	27	27	27
Income: \$100k-200k (%)	30	30	30
Income: \$200k+ (%)	16	16	16

Note: Table 1 reports summary statistics for our baseline sample of consumers, described in Section 2, as of January 2025, July 2025, and January 2026. The age and income shares sum to less than 100% since some consumers may have missing demographics.

Table 2: Switching Between OpenAI and Anthropic Around Model Improvements

	Switch Rate (%)				
	(1)	(2)	(3)	(4)	(5)
Model improved in t	0.5136** (0.2521)			0.6903** (0.2757)	
Model improved in $t - 1$		0.4144 (0.2688)		0.6287** (0.2945)	
Model improved in $t - 2$			-0.1079 (0.2212)	0.1490 (0.2115)	
Improvement amount in t and $t - 1$					0.0950* (0.0529)
Provider FE	Yes	Yes	Yes	Yes	Yes
Observations	54	54	54	54	54
R^2 (within)	0.0857	0.0539	0.0036	0.1919	0.0713
Mean dep. var.	2.7807	2.7807	2.7807	2.7807	2.7807

Note: Table 2 shows regressions of the switching rate on measures of model improvement from the Epoch Capabilities Index (ECI). The dependent variable is the share of consumers who subscribed exclusively to the other provider in the previous month and switch to or add the given provider in the current month. The unit of observation is the provider-month for OpenAI and Anthropic. "Model improved" is an indicator for whether the provider released a model that improved its frontier ECI score. "Improvement amount" is the size of the ECI score change. All specifications include provider fixed effects. Heteroskedasticity robust standard errors are reported in parentheses. Significance levels: *($p < 0.1$), **($p < 0.05$), ***($p < 0.01$).

Table 3: Estimated Consumer Preferences

(a) Mean Preferences					(b) Price and Quality Coefficient			
	OpenAI	Anthropic	GoogleAI	xAI		OLS	IV2SLS	
Q2 2024	-16.0 (0.004)	-18.0 (0.026)	-18.8 (0.020)	-20.8 (0.045)	Price ($-\alpha$)	-0.179 (0.049)	-0.222 (0.039)	
Q3 2024	-15.7 (0.004)	-17.4 (0.024)	-18.1 (0.018)	-18.4 (0.037)	Quality (β)	0.160 (0.034)	0.167 (0.033)	
Q4 2024	-15.1 (0.004)	-17.0 (0.023)	-17.8 (0.018)	-18.0 (0.037)				
Q1 2025	-14.7 (0.003)	-16.6 (0.023)	-17.5 (0.018)	-18.0 (0.037)	(c) Preference Correlation			
Q2 2025	-13.8 (0.003)	-16.4 (0.022)	-17.4 (0.018)	-17.9 (0.037)	OpenAI	1.00	0.889	0.545
Q3 2025	-13.2 (0.003)	-16.0 (0.022)	-17.1 (0.017)	-17.7 (0.037)	Anthropic	0.889	1.00	0.758
Q4 2025	-12.7 (0.003)	-15.6 (0.021)	-16.7 (0.017)	-17.3 (0.037)	GoogleAI	0.545	0.758	1.000
Q1 2026	-12.1 (0.003)	-14.3 (0.020)	-16.2 (0.017)	-17.0 (0.037)	xAI	0.562	0.781	0.999
Std. Dev.	6.00	5.54	5.52	5.66				
(d) Complementarity Parameters								
	Anth. $\times$ Goog.	Anth. $\times$ Open.	Anth. $\times$ xAI	Goog. $\times$ Open.	Goog. $\times$ xAI	Open. $\times$ xAI		
Q2 2024	-9.13 (31.4)	-4.10 (0.031)	-12.9 (140)	1.49 (0.046)	-8.64 (5.28)	3.22 (0.067)		
Q3 2024	0.516 (0.141)	-4.55 (0.028)	-2.71 (0.525)	1.22 (0.033)	-6.80 (4.94)	2.50 (0.026)		
Q4 2024	-1.57 (0.275)	-4.99 (0.025)	-2.85 (0.506)	0.685 (0.029)	-10.7 (5.01)	1.90 (0.023)		
Q1 2025	-5.50 (1.51)	-5.28 (0.024)	-2.66 (0.499)	0.278 (0.027)	-9.07 (4.95)	1.82 (0.021)		
Q2 2025	0.773 (0.063)	-5.58 (0.022)	-2.04 (0.492)	-0.040 (0.023)	-7.13 (4.94)	1.11 (0.020)		
Q3 2025	0.642 (0.053)	-5.81 (0.021)	-1.75 (0.489)	-0.199 (0.018)	-6.28 (4.94)	0.865 (0.018)		
Q4 2025	0.251 (0.047)	-6.14 (0.020)	-1.90 (0.488)	-0.494 (0.016)	-6.75 (4.94)	0.374 (0.017)		
Q1 2026	-0.156 (0.033)	-6.44 (0.018)	-2.14 (0.487)	-0.905 (0.015)	-6.61 (4.94)	-0.043 (0.017)		
Std. Dev.	2.64	5.08	0.043	1.13	2.42	0.710		

Note: Table 3 reports estimated preference parameters from the demand model, with standard errors in parentheses. Panel (a) reports mean preference parameters and the standard deviations of preference heterogeneity. Panel (b) reports estimates of price sensitivity and the valuation of model quality. Panel (c) reports correlations in preference heterogeneity across providers. Panel (d) reports the mean and standard deviation of complementarities across AI providers over time. For the standard errors in panels (a) and (d), we compute the Hessian matrix of the negative simulated log-likelihood at the estimated parameter vector using JAX automatic differentiation, invert it to obtain the asymptotic covariance matrix, and take the square root of the diagonal elements to recover parameter-level standard errors.

Table 4: Estimated Provider Markups

	Static	Dynamic
<i>Markups (\$/month)</i>		
OpenAI	\$11.8	\$4.07
Anthropic	\$8.62	\$6.96
GoogleAI	\$7.45	\$6.67
xAI	\$5.77	\$5.34
<i>Lerner Index</i>		
OpenAI	59.0%	20.3%
Anthropic	43.1%	34.8%
GoogleAI	37.3%	33.3%
xAI	19.2%	17.8%

Note: Table 4 reports estimated markups under both static and dynamic supply-sides.

Table 5: Price and Share Effects in Counterfactuals

	Baseline	$\Gamma = 0$	$ \Gamma $	Uncorr. Prefs.
<i>Prices (\$/month)</i>				
OpenAI	\$20.0	\$21.7	\$26.7	\$20.4
Anthropic	\$20.0	\$26.2	\$32.7	\$20.8
GoogleAI	\$20.0	\$21.3	\$27.1	\$21.0
xAI	\$30.0	\$32.5	\$39.2	\$32.6
<i>Shares</i>				
OpenAI	2.60%	2.66%	3.02%	2.58%
Anthropic	0.503%	0.962%	2.33%	0.648%
GoogleAI	0.242%	0.239%	0.352%	0.236%
xAI	0.124%	0.146%	0.335%	0.161%

Note: Table 5 reports estimated prices and market shares in three counterfactuals with an optimizing supply-side in 2026Q1. The first counterfactual sets net substitutability, Γ , equal to zero. The second counterfactual sets net substitutability to be positive, $|\Gamma|$. The third counterfactual sets the correlation of preferences to be zero (i.e., panel (c) in Table 3). All counterfactuals hold other parameters fixed.

Appendix

A Additional Details on Empirical Analyses

A.1 Identifying AI Transactions

We identify transactions from AI providers using information on merchant names and transaction sizes. For OpenAI, Anthropic, and xAI we use only information on merchant names. We match on a variety of provider and product name strings (e.g., both "OpenAI" and "ChatGPT") and manually validate the matches to exclude false positives and negatives. For Google, identifying AI transactions is more challenging because the company processes many non-AI payments and bundles its Gemini AI service with other products through Google One, where only higher-priced tiers include AI features. We restrict to transactions whose merchant descriptors contain identifiers specific to Google's AI subscription tiers (including Google One), impose a minimum transaction amount consistent with the pricing of these tiers, and require that flagged consumers also appear in at least one adjacent month, to exclude annual purchases of lower-tier Google One plans that do not include AI features.

B Additional Details on Model

B.1 Estimation Details

This appendix provides more detail on how we estimate the demand side of the model in Section 4.

Data. We use panel data on adoption across 9 quarters from 2024Q1-2026Q1. We drop consumers that use more than two bundles in any period or have less than 20 transactions in any period.

Simulated Log-Likelihood. We simulate the log-likelihood by first generating $R = 100$ draws from scrambled Halton sequences. Each draw is dimension $4 + 6 = 10$ to integrate over θ and γ . These Halton draws are generated once. To recover θ^r , we transform the Halton draws to standard normal draws and multiply by L'_θ , where L_θ is the lower-triangular Cholesky factor of Σ_θ . For γ^r , there is no covariance structure so each standard-normal draw just gets scaled by the square root of the corresponding entry in Σ_γ . Given θ^r and γ^r , we compute $s_{ict}(\theta^r, \gamma^r)^{d_{ict}}$ and the likelihood in equation (19).

Optimizer. To find the set of parameters that minimize the negative log likelihood, we start by searching over an initial set of starting points. The first starting value of δ_{jt} is the difference

in log shares for product j in period t relative to the outside option in period t . Γ_{ct} is set to -5 for all pairs to begin with. L_θ is set to a matrix of all ones on the lower diagonal. Σ_γ is zero for all pairs. We then use Latin hypercube sampling around these initial values to generate 1,000 additional possible starting points. We evaluate the log-likelihood at each of these 1,000 starting points, and pick the five with the lowest negative log likelihood to “polish.” For each of these five starting points, we use L-BFGS-B with a functional tolerance of 10^{-9} and gradient tolerance of 10^{-5} . Among the five sets of polished parameters, we select the set of parameters that can achieve the lowest negative log likelihood.

B.2 Calibrating the Dynamic Supply Side

We estimate the dynamic supply side in two steps. First, we estimate the evolution of product quality using a simple linear regression. Second, we calibrate each provider’s marginal cost, κ_j , using the dynamic first-order conditions that rationalize observed pricing decisions.

Evolution of Product Quality. We estimate the evolution of product quality by ordinary least squares:

$$\omega_{jt} = \rho_0 + \rho\omega_{j,t-1} + \lambda s_{j,t-1} + \phi_t + \eta_{jt}. \quad (\text{B.1})$$

The parameter ρ captures persistence in product quality, while λ captures the extent to which a larger user base contributes to future product quality through learning by doing. The time fixed effects, ϕ_t , capture changes in the state of AI technology that are common across providers, while η_{jt} captures idiosyncratic innovations in provider j ’s model quality.

Supply Calibration. We calibrate the marginal costs of each provider using the steady-state pricing conditions implied by the dynamic model. For tractability, we treat the observed market as a steady state when calibrating the supply side. This allows us to characterize firms’ continuation values using the steady-state envelope conditions rather than numerically solving the full dynamic optimization problem. Given the estimated demand parameters, the observed vector of prices $\mathbf{p}$, and the observed vector of model qualities $\boldsymbol{\omega}$, we construct the matrices

$$D \equiv \frac{\partial \mathbf{s}}{\partial \mathbf{p}}, \quad H \equiv \frac{\partial \mathbf{s}}{\partial \boldsymbol{\omega}}, \quad (\text{B.2})$$

and the Jacobian of the quality transition,

$$A \equiv \rho I + \lambda H. \quad (\text{B.3})$$

The steady-state envelope conditions imply that the matrix of marginal values of future model quality is

$$\mathcal{V} = \text{diag}(\mathbf{p} - \boldsymbol{\kappa})H(I - \rho A)^{-1}. \quad (\text{B.4})$$

Substituting this expression into the firms' dynamic first-order conditions yields a closed-form expression for marginal costs:

$$\kappa = \mathbf{p} + [\text{diag}(D) + \varrho\lambda \text{diag}(H(I - \varrho A)^{-1}D)]^{-1} \mathbf{s}. \quad (\text{B.5})$$

Thus, conditional on the estimated demand and quality-transition parameters, marginal costs are recovered directly from observed prices, market shares, and model qualities. When $\lambda = 0$, the dynamic learning channel disappears and equation (B.5) reduces to the marginal-cost recovery implied by the static pricing model.

B.3 Solving for the Counterfactual Dynamic Steady-State Equilibrium

We analyze counterfactuals by comparing steady-state equilibria under alternative market structures and model parameters. We hold the estimated demand parameters, quality-transition parameters, and calibrated marginal costs fixed unless they are explicitly changed as part of the counterfactual. For each counterfactual environment, we jointly solve for the vector of equilibrium prices, $\mathbf{p}^*$, and model qualities, ω^* .

A counterfactual steady state must satisfy two sets of conditions. First, model quality must be stationary. We set the common component of technological progress to zero, $\phi = 0$, so that differences in steady-state model quality across counterfactuals arise entirely from differences in provider-specific learning. From equation (B.1), steady-state quality therefore satisfies

$$\omega^* = \rho_0 + \rho\omega^* + \lambda s(\mathbf{p}^*, \omega^*), \quad (\text{B.6})$$

or equivalently,

$$\omega^* = \frac{\lambda s(\mathbf{p}^*, \omega^*) + \rho_0}{1 - \rho}. \quad (\text{B.7})$$

Thus, a provider's steady-state model quality is directly related to the size of its user base. Counterfactual changes in market structure affect prices and the allocation of consumers across providers, which in turn determine the amount of learning accumulated by each provider and its long-run model quality.

Second, prices must satisfy each provider's dynamic first-order condition. At a candidate $(\mathbf{p}^*, \omega^*)$, we compute

$$D^* \equiv \frac{\partial \mathbf{s}}{\partial \mathbf{p}} \Big|_{(\mathbf{p}^*, \omega^*)}, \quad H^* \equiv \frac{\partial \mathbf{s}}{\partial \omega} \Big|_{(\mathbf{p}^*, \omega^*)}, \quad (\text{B.8})$$

and

$$A^* = \rho I + \lambda H^*. \quad (\text{B.9})$$

The corresponding matrix of marginal continuation values is

$$\mathcal{V}^* = \text{diag}(\mathbf{p}^* - \boldsymbol{\kappa})H^*(I - \varrho A^*)^{-1}. \quad (\text{B.10})$$

Equilibrium prices therefore satisfy


$$\mathbf{s}^* + \text{diag}(D^*)(\mathbf{p}^* - \boldsymbol{\kappa}) + \varrho\lambda \text{diag}(\mathcal{V}^*D^*) = \mathbf{0}, \quad (\text{B.11})$$

where $\mathbf{s}^* = \mathbf{s}(\mathbf{p}^*, \boldsymbol{\omega}^*)$.

We solve equations (B.6) and (B.11) jointly for $(\mathbf{p}^*, \boldsymbol{\omega}^*)$. We use the resulting counterfactual equilibrium to compare prices, adoption, model quality, profits, and consumer welfare across alternative market structures.


C Additional Figures and Tables

Figure C.1: Geography of Paid AI Adoption in Massachusetts



Note: Figure C.1 maps the share of consumers paying for AI across ZIP codes in Massachusetts as of April 2026.


Figure C.2: Geography of Paid AI Adoption in New York City




Note: Figure C.2 maps the share of consumers paying for AI across ZIP codes in New York City as of April 2026.

Figure C.3: Adoption Across Products by Commuting Zone

(a) OpenAI versus Anthropic




(b) OpenAI versus xAI




Note: Figure C.3 compares paid AI adoption across commuting zones as of April 2026. Each point represents a commuting zone. Panel (a) plots the number of paid OpenAI users per 100,000 consumers against the number of paid Anthropic users per 100,000 consumers. Panel (b) similarly compares paid OpenAI and xAI users. Dashed lines show linear fits.

Figure C.4: Other Products Around the Release of ChatGPT



Note: Figure C.4 shows changes in the number of consumers paying for selected products around the release of ChatGPT. The horizontal axis reports months relative to the release of ChatGPT, with month zero corresponding to the release month. The vertical axis reports the percentage change in paid users. The vertical dashed line denotes the release of ChatGPT, and the horizontal dashed line denotes zero. Panels (a)–(d) show results for Chegg, Midjourney, Grammarly, and Adobe, respectively.

Figure C.5: Spending Around OpenAI Product Adoption



Note: Figure C.5 shows spending in the 12 months before and after a consumer's first purchase of an OpenAI product. Each panel shows spending in a different Merchant Category Code (MCC) among OpenAI adopters and matched non-adopters. We match each adopter to a random non-adopter in the same state-by-age-group-by-income-group cell. Month zero denotes the month of the adopter's first OpenAI purchase. Non-adopters are assigned a placebo adoption date in the same month as their matched adopter. Both groups are restricted to a balanced panel of individuals with at least 10 transactions in every month from January 2022 through April 2026. Computer software spending excludes purchases from OpenAI but includes purchases of other AI products.

Table C.1: Determinants of AI Adoption

	(1) Anthropic	(2) Google	(3) OpenAI	(4) xAI	(5) Midjourney	(6) QuillBot	(7) Cursor
<i>Age (base: Young, <35)</i>							
Working age (35–54)	-0.0041*** (0.0003)	0.0045*** (0.0002)	-0.0510*** (0.0010)	-0.0009*** (0.0002)	-0.0098*** (0.0003)	-0.0014*** (0.0001)	-0.0026*** (0.0001)
Pre-retirement (55–65)	-0.0167*** (0.0003)	-0.0044*** (0.0002)	-0.2031*** (0.0013)	-0.0037*** (0.0002)	-0.0256*** (0.0003)	-0.0032*** (0.0001)	-0.0047*** (0.0001)
Retired (65+)	-0.0299*** (0.0002)	-0.0067*** (0.0002)	-0.4076*** (0.0009)	-0.0055*** (0.0002)	-0.0362*** (0.0003)	-0.0087*** (0.0001)	-0.0072*** (0.0001)
<i>Income (base: Low, <\$50k)</i>							
Moderate (\$50k-100k)	0.0028*** (0.0002)	-0.0020*** (0.0001)	0.0344*** (0.0007)	-0.0003* (0.0001)	0.0017*** (0.0002)	0.0004*** (0.0001)	0.0004*** (0.0001)
High (\$100k-200k)	0.0009*** (0.0002)	-0.0033*** (0.0001)	0.0501*** (0.0008)	-0.0006*** (0.0001)	0.0005** (0.0002)	0.0011*** (0.0001)	-0.0002** (0.0001)
Elite (\$200k+)	0.0149*** (0.0003)	-0.0037*** (0.0002)	0.2004*** (0.0014)	0.0032*** (0.0002)	0.0006** (0.0002)	0.0040*** (0.0001)	0.0023*** (0.0001)
<i>Zip-code demographics</i>							
College share	0.2027*** (0.0008)	0.0386*** (0.0004)	1.5467*** (0.0031)	0.0563*** (0.0004)	0.1418*** (0.0006)	0.0129*** (0.0002)	0.0332*** (0.0003)
White share	-0.0364*** (0.0013)	-0.0212*** (0.0010)	-0.3023*** (0.0054)	-0.0174*** (0.0008)	-0.0883*** (0.0011)	-0.0039*** (0.0007)	-0.0014** (0.0005)
Asian share	0.1187*** (0.0034)	0.0120*** (0.0016)	0.8462*** (0.0167)	0.0035* (0.0014)	-0.0366*** (0.0019)	0.0235*** (0.0010)	0.0512*** (0.0012)
Black share	-0.0215*** (0.0013)	-0.0093*** (0.0010)	-0.1326*** (0.0056)	-0.0279*** (0.0008)	-0.0711*** (0.0012)	0.0199*** (0.0007)	-0.0021*** (0.0005)
Hispanic share	-0.0037*** (0.0010)	-0.0075*** (0.0008)	0.0962*** (0.0046)	-0.0120*** (0.0007)	-0.0431*** (0.0009)	0.0135*** (0.0006)	0.0018*** (0.0004)
Observations	21,238,509	21,238,509	21,238,509	21,238,509	21,238,509	21,238,509	21,238,509
R^2 (within)	0.0260	0.0019	0.1148	0.0031	0.0115	0.0027	0.0072
Month FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: Table C.1 reports estimates from a regression of AI adoption on consumer characteristics, conditional on month fixed effects. Each observation is a ZIP code $\times$ age group $\times$ income group $\times$ month cell, weighted by the number of individuals in the cell. The outcome variable is the share of individuals with at least one paid transaction to each provider, multiplied by 100 (percentage points). Heteroskedasticity-robust standard errors are reported in parentheses. Significance levels: *($p < 0.1$), **($p < 0.05$), ***($p < 0.01$).