AI and the Market for Expert Services: Evidence from Financial Advisers

Federico Mainardi
Columbia Business School

Dongwoo Noh
HKUST

Sangmin S. Oh
Columbia Business School

Sean Seunghun Shin*
KAIST

August 2026

Preliminary. Please Do Not Distribute Without Explicit Permission.

ABSTRACT

AI is widely predicted to commoditize expert advice and displace the human experts who provide it. We examine this prediction in the market for financial advice, where households delegate wealth to advisers for a bundle of technical portfolio guidance and relationship-specific trust. On the demand side, we show in customer disputes that the advice households contest is overwhelmingly technical rather than relational, and that AI can both audit it and produce it. On the supply side, advisers adopt AI on this same technical margin, albeit slowly. We formalize this interaction in a model in which AI lowers the cost of producing, checking, and replacing technical advice, but not the cost of building trust. The model shows that, with the advent of AI, (i) the fee response is ambiguous, (ii) fees and advised assets cannot both decline, and (iii) the profit response to AI is more favorable for advisers who produce trust cheaply. AI may therefore commoditize technical advice without displacing advisers themselves, shifting advantage toward those who sustain trusted relationships both in financial advice and in expert services more broadly.

Keywords: artificial intelligence; financial advisers; expert services; credence goods

JEL Codes: G20, G23, G28, D82, L11, O33

*First Draft: August 2026. For comments and helpful feedback, we thank Maxim Alekseev, Xuelin Li and Harry Mamaysky. Mainardi: fm2913@columbia.edu. Oh: oh@gsb.columbia.edu. Noh: noh@ust.hk. Shin: sean.shin@kaist.ac.kr.

Generative artificial intelligence (AI) has diffused rapidly, with ChatGPT reaching an estimated one hundred million users within two months of its release. This mass adoption has prompted predictions that systems capable of expert-level reasoning at very low cost will displace much knowledge work within years rather than decades.[1] The market for financial advice provides a natural and economically important setting in which to examine this prediction. Registered investment advisers (RIAs) manage trillions of dollars for U.S. households, and AI can already produce personalized saving and portfolio recommendations (Choukhmane et al., 2026). A recent market episode illustrates the perceived competitive threat: the announcement of an AI-powered adviser tool coincided with declines of nearly 10 percent in the share prices of publicly traded advisory firms (Figure 1).

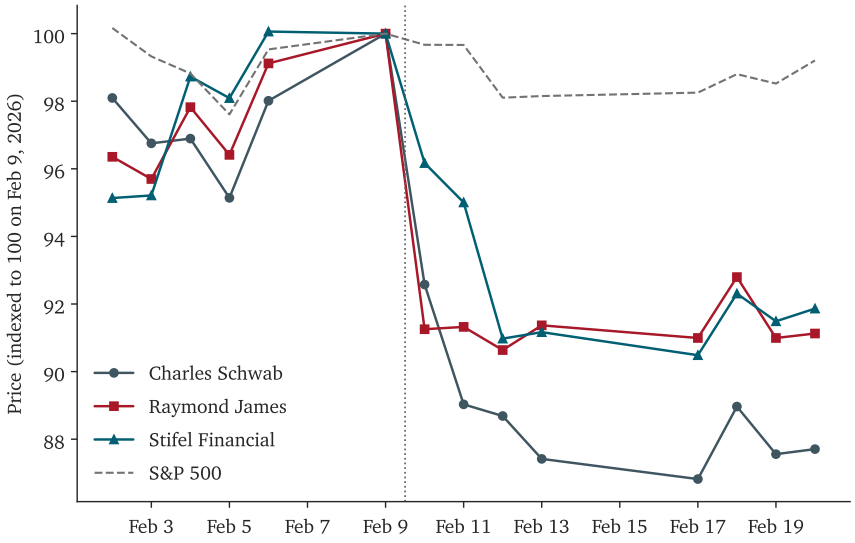


Figure 1. **Stock-Market Reaction to the Launch of an AI-Powered Adviser.** This figure plots daily closing prices of three publicly traded advisory and brokerage firms and the S&P 500 index, indexed to 100 on February 9, 2026, the trading day before Altruist announced Hazel, an AI-powered tax-planning tool for advisers. The dotted vertical line marks the announcement. (Data Source: Compustat)

In this paper, we study how generative AI reshapes the market for financial advice. We document how advisers use AI, assess whether clients can use the same technology to evaluate the advice they receive and generate alternatives, and use these findings to develop a model of fees, delegation, and adviser profits.

Two features of financial advice are central to our analysis. First, financial advice is a credence good: clients cannot readily assess its quality even after acting on it (Dulleck and Kerschbamer, 2006). The quality of a recommendation depends on how well it fits the client's

[1]For example, in 2016, Geoffrey Hinton declared that "people should stop training radiologists now." More recently, GPT-4 achieved a passing score on a simulated Uniform Bar Examination (Katz et al., 2024).

circumstances, but realized returns are noisy—a well-suited portfolio can lose money in a bad market, while a concentrated bet can perform well in a good one. Short-term performance therefore reveals little about the quality of the advice. Second, financial advice bundles *technical guidance*—security selection, asset allocation, and monitoring—with a relational component whose value is tied to the human provider (Garicano, Li, and Wu, 2026; Imas, 2026). Relationship-specific trust, in turn, helps households act on that guidance and sustain delegation (Gennaioli, Shleifer, and Vishny, 2015; Greig et al., 2025). Combined with clients' direct access to generative AI, these features make the technology a two-sided shock rather than a conventional supply-side automation technology. Advisers can use AI to lower the cost of producing technical guidance, while households can use the same tools to evaluate that guidance and obtain alternatives. The relational component limits direct substitution away from advisers, but it does not fully insulate them: by making technical guidance easier to evaluate and replace, AI strengthens households' outside option. The market consequences therefore depend on how these supply- and demand-side forces interact.

We begin on the demand side, using customer disputes to identify which advisory services clients most often contest—and where an independent evaluation may therefore have been especially valuable. We assemble nearly 24,000 disputes from BrokerCheck and the Investment Adviser Public Disclosure (IAPD) database, public registries covering registered brokers and investment advisers. Each record contains a household's allegation in narrative form, allowing us to identify the service at the center of the dispute. We classify these narratives using a taxonomy of advisory services and find that nearly three quarters concern technical services, such as security selection. Relational services central to delegation account for about one percent.[2] These shares do not characterize all advice households receive. Rather, they show that among relationships serious enough to generate formal disputes, the contested margin is overwhelmingly technical. We next ask whether AI can actually audit that margin and provide an alternative.

We evaluate both capabilities in a blinded design that adapts the audit-study logic of Mullainathan, Noeth, and Schoar (2012) to AI. The main empirical challenge is that dispute records mix information available when advice was given with ex-post accounts of what followed. We therefore use one LLM to reconstruct the client's situation when advice was sought and the adviser's recommendation, preserving contemporaneous claims but removing ex-post characterizations. We also mask details that could reveal the dispute or its outcome. An LLM "judge" from a separate model family then scores each recommendation from 1 (clearly inappropriate) to 5 (well-suited). Each case is evaluated independently, and the judge does not know that a

[2]The remaining 26.4 percent of disputes fall in categories classified as mixed rather than technical or relational. Nearly all of it is risk and insurance planning (23.9 percent), chiefly allegations over annuity and insurance sales.

complaint was filed.

Using only this reconstructed ex-ante information, the judge rates disputed human advice poorly. The mean score is 2.1, and three quarters of recommendations receive a 1 or 2. Because the judge observes neither outcomes nor dispute status, these results indicate that problems in advice that later became the subject of formal complaints were often detectable from information available when the recommendation was made. AI can therefore provide households with an outside check on advice quality.

Auditing alone, however, does not give households a full outside option; AI must also be able to generate an alternative. For each case, a separate model observes only the reconstructed client situation—not the human recommendation or the existence of a dispute—and produces its own recommendation. We standardize the human and AI recommendations to conceal authorship and present them to the judge in randomized order without source labels. AI-generated advice receives a mean score of 4.3, and the judge prefers it to the corresponding human advice in 95 percent of cases. Together, the exercises provide evidence that AI can both audit human advice and offer a credible alternative on the technical margin.

Second, on the supply side, we show that advisers adopt AI on the same technical margin, albeit slowly. To measure adoption, we use Form ADV Part 2 brochures, the narrative filings in which every registered adviser describes its services in its own words. The brochures are mandatory and must be updated at least annually and promptly upon material change, so a firm that incorporates AI into its services has the regulatory incentive to disclose it. To measure AI adoption, we first flag AI-related language with a deterministic keyword screen. An LLM then reads every flagged brochure and separates firms that affirmatively use AI from risk-disclosure boilerplate or investment themes. Every label is supported by a verbatim excerpt from the filing, and we manually verify the classifications against these excerpts and the underlying filings.

Our main finding is that adoption of AI by advisers is real but still early, and that the use is concentrated in the production of technical advice. Disclosed use is essentially absent before 2021, and the number of advisers disclosing use rises from 44 in 2021 to 350 in 2025. Even so, these counts are small relative to the more than ten thousand advisers in our sample. Moreover, the use of AI sits on the technical margin: 72 percent of use brochures place AI in investment-process tasks—screening securities, constructing portfolios, generating recommendations—and 50 percent in research. Firms rarely sell AI as a separate product: only 5.3 percent of use brochures mention AI in a fee section, and 6.5 percent describe it as a separately branded program or service. These findings suggest that advisers absorb AI slowly, as an input into the production of technical advice, not as a new product line with its own price.

We also find that adoption varies sharply with the clientele an adviser serves. If institution-facing advisers specialize more in selling technical advice rather than relationship-specific trust, they should have greater use for a technology that makes such advice cheaper to produce. Adoption indeed declines in the household share of an adviser's client assets, from 8.9 percent among purely institutional advisers to 2.3 percent among the most household-facing, with the intermediate groups falling in between. Overall, the cross-section of adoption is consistent with the interpretation that AI lowers the cost of producing technical advice, and the advisers who specialize in that production have adopted it first.

Taken together, the demand- and supply-side evidence gives the same technology two roles. It lowers the adviser's cost of producing technical advice while strengthening the household's outside option, so the two forces pull the price of advice in opposite directions. With adoption still early, where fees, the size of the advisory sector, and adviser profits eventually settle is an equilibrium question. We therefore develop a model in which the technology makes technical advice cheaper for the adviser to produce and cheaper for the household to check and replace, while leaving trust beyond its reach on either side. Comparative statics in the price of machine intelligence then trace how each outcome responds.

We build on the model of Gennaioli, Shleifer, and Vishny (2015), in which households delegate their portfolios to advisers they trust, and trust lowers the cost of bearing risk: the same risky position feels less fraught when a trusted adviser recommends it. To this framework we add the technical margin that AI reaches. The adviser chooses a fee and the levels of two inputs: a technical input, which raises the expected return the client earns and whose cost falls with the price of machine intelligence, and a relational input, which deepens the relationship and whose cost AI leaves unchanged. The household, in turn, decides whether to delegate its portfolio at the posted fee or to invest on its own. Using AI sharpens what the household knows about the advice it receives and raises the return the household can earn without an adviser. We represent the arrival of AI as a fall in the price of machine intelligence, which reaches the adviser's production cost and the household's use of the technology alike.

The first result is that the fee response to AI is ambiguous. In equilibrium, the fee decomposes into a return rent, the adviser's technical edge over the household's outside option, and a trust premium attached to the relationship. A fall in the price of machine intelligence moves both components of the fee. Cheaper production raises the return the adviser delivers, which raises the rent she can extract. Furthermore, a larger book raises the marginal value of the relationship, so the adviser uses more of the relational input, and the trust premium increases as well. On the other hand, the household's better outside option compresses what the adviser can charge. Therefore, the common intuition that cheaper production must lower prices holds

only when the outside-option force dominates.

The second result links fees to the size of the advisory sector, measured by advised assets. Every non-fee channel through which the technology operates expands the amount households are willing to delegate: better returns raise the value of being advised, more trust lowers the cost of bearing risk, and sharper auditing reduces the uncertainty that keeps households out. A lower fee expands delegation further. Advised assets can therefore shrink only if fees rise. In other words, the popular scenario that AI compresses fees and that also hollows out the advisory sector cannot both be true in the model.

The third result concerns which advisers gain or lose. Advisers can differ in the cost of producing trust: a firm that produces trust cheaply builds its business around relationships, while one that cannot do so tilts toward technical services. A comparative static in this cost shows that following a fall in the price of machine intelligence, the profit response is weaker for advisers with a high cost of producing trust than for advisers who produce it cheaply. We call this prediction *trust cover*: AI treats advisers whose business rests on the relationship more favorably than those whose business rests on technical production.

Taken together, our results suggest that AI is unlikely to displace human financial advisers as a whole, but it is likely to reprice what they sell. The technology commoditizes the component of advice that can be checked and reproduced, while the trust that makes households act on advice remains outside its reach. Whether an adviser gains or loses therefore depends not only on exposure to automation but also on the composition of what the adviser sells. This conclusion likely extends beyond financial advice. Law and medicine share a similar structure: an expert sells checkable technical claims bundled with judgment, accountability, and a relationship (Autor and Thompson, 2025; Garicano, Li, and Wu, 2026). Our results suggest that a technology that reaches the auditable component of expert work but not the relational one does not eliminate the expert but instead changes what the expert is paid for.

Related Literature

This paper first contributes to the household finance literature on financial advisers. The conceptual benchmark is the "money doctors" view of advice as a bundle of information and trust (Gennaioli, Shleifer, and Vishny, 2015). We extend this logic by studying a technology that changes the value of one part of the bundle: AI makes technical portfolio guidance easier for households to check and self-produce, while leaving relationship-specific trust as a distinct source of delegation value. Empirically, the paper relates to evidence on how advisers shape household portfolios (Foerster et al., 2017; Chalmers and Reuter, 2020) and how agency fric-

tions appear in advisory relationships (Egan, Matvos, and Seru, 2019). It also relates to work on the quality of advice that households receive (Mullainathan, Noeth, and Schoar, 2012) and the fees advisers charge (Mainardi and Oh, 2026).[3] Our primary contribution to this literature is to study how a new information technology reshapes the market for human financial advice.

Our view of financial advice is also related to the literature on credence goods. In classic credence-good settings, such as medicine, law, and expert repair, consumers rely on specialists because they cannot fully assess the quality or necessity of the service they receive (Darby and Karni, 1973; Dulleck and Kerschbamer, 2006). Financial advice has this feature on its technical margin: households often cannot know whether a portfolio recommendation is appropriate at the time advice is given. AI introduces a new dimension to this problem as it changes which parts of expert advice can be checked and self-produced by consumers. The same logic may apply beyond finance to other expert-service markets in which AI makes some technical claims more contestable while leaving relationship-specific delegation valuable.

Second, the paper relates to the task-based view of automation, in which technology acts on individual tasks rather than on whole occupations. Because jobs bundle many tasks, automating one component can either erode or raise the value of the remaining work (Autor, Levy, and Murnane, 2003; Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2019). Recent work applies this task-based logic to AI, measuring exposure to large language models and estimating their effects on productivity, task allocation, and workflow design (Eloundou et al., 2023; Brynjolfsson, Li, and Raymond, 2023; Dell'Acqua et al., 2023; Chatterji et al., 2026). We bring this logic to financial advice. Advisers perform such a bundle, with analytical tasks on one side and relationship-building tasks on the other. AI changes the value of this bundle by standardizing some technical tasks while increasing the relative importance of the remaining relationship-specific work.

Our work also pertains to the growing literature studying the implications of AI for the financial sector. Recent work studies whether AI can perform various financial tasks (Lopez-Lira and Tang, 2023; Wu et al., 2023; Kim, Muhn, and Nikolaev, 2024) and how access to these tools shapes retail investment decisions (Kim et al., 2024; Blankespoor, Croom, and Grant, 2026), though AI-generated advice can carry biases of its own (Fedyk et al., 2025). Closest to our setting, recent work evaluates the quality of AI-generated personal financial advice and asks whether LLM recommendations align with benchmark economic models (Choukhmane et al., 2026).[4] We differ by shifting the focus from whether AI can perform financial tasks to how its availability changes the market for human advice. Our contribution is to treat these two

[3] See also Reuter and Schoar (2024) for a review of the demand- and supply-side constraints in this market.
[4] See Mo and Ouyang (2025) for a review of AI in financial economics.

margins jointly and trace their equilibrium consequences for the advisory industry.

Third, the paper also relates to research on earlier waves of technology in retail finance, most prominently robo-advising. Robo-advising studies show how algorithmic portfolio tools can standardize investment advice, reduce behavioral mistakes, and expand access to low-cost portfolio management (D'Acunto, Prabhala, and Rossi, 2019; Rossi and Utkus, 2020; D'Acunto and Rossi, 2023). A broader fintech literature studies how digital platforms, alternative data, and machine learning change screening, pricing, access, and market structure in consumer finance (Buchak et al., 2018; Fuster et al., 2019; Berg et al., 2020; Bartlett et al., 2022). AI differs from these technologies in two ways. First, it is not only a supply-side automation technology; it is also available to households as a tool for evaluating and replacing parts of advice. Second, whereas robo-advice largely automates portfolio allocation, AI enters a broader advice relationship that bundles technical guidance with communication, trust, and delegation. In this sense, our results extend this literature from the automation of a single task to a technology that reaches both sides of the adviser-client relationship.

1 Institutional Background and Data

We first characterize the nature of financial advice and describe the market for financial advisers. We then describe the data used in the empirical analysis.

1.1 The Advisory Industry

Financial advice in the United States is primarily provided by registered investment advisers, firms that owe their clients a fiduciary duty and register with the SEC or, below an asset threshold, with state regulators. As of 2025, over 16,000 SEC-registered advisers manage nearly $180 trillion in client assets (Investment Adviser Association, 2026). Like other expert services such as law and medicine, the industry sells professional judgment that clients cannot easily evaluate, under an accountability standard that exists precisely because clients cannot judge quality directly.

The product these firms sell is a bundle of technical and relational advice. Technical advice consists of portfolio guidance, including asset allocations, product selection, and liquidity management. Relational advice consists of relationship-specific services: advisers provide reassurance to clients during market stress and communicate in ways that sustain trust. Liability is the extreme relational margin: an accountable professional stands behind the recommenda-

tion and absorbs part of the responsibility for the advice process.

Our unit of analysis is the advisory firm rather than the individual financial adviser. The main outcomes of interest, such as advisory fees and the adoption of AI, are determined at the firm level, and the regulatory filings that record them are also filed by firms. Within this population, we focus on advisers that serve households, defined as firms reporting positive regulatory assets under management from individual clients, whether retail or high net worth, in at least one year of our sample period, similar in spirit to the household-facing advisers studied by Mamaysky and Zhang (2025).

1.2 Data

We combine several regulatory data sources that observe different parts of the advisory relationship. First, to study which advice households contest and whether AI can evaluate it (Section 2), we use customer-dispute records from FINRA BrokerCheck and the SEC's Investment Adviser Public Disclosure (IAPD) database. FINRA maintains BrokerCheck for registered brokers and brokerage firms, while the SEC maintains IAPD for registered investment advisers and their representatives. Together, the two registries provide public disclosure coverage across the principal regulatory channels through which financial professionals serve U.S. clients.[5] We use the underlying allegation narratives to identify which advisory service each dispute concerns and whether the disputed advice can be evaluated by AI.

To study advisers' own adoption of AI (Section 3), we use Form ADV filings. Every registered adviser must file Form ADV annually. The form has two parts: Part 1 is a structured questionnaire filed electronically, reporting assets under management, client composition, ownership, and compensation arrangements; Part 2 is a narrative "brochure" delivered to clients that describes, in free text, the adviser's services, fees, and conflicts of interest. From Part 1 we draw firm characteristics, including regulatory assets under management (RAUM) by client type in Item 5.D, which we use to measure each adviser's clientele composition. The Part 2 brochures provide the firm's own description of its advisory services and form the basis of the AI-adoption measure in Section 3.

Our sample covers years 2018–2025. To construct the sample, we retain the registered advisers that report positive regulatory assets from either retail clients or high-net-worth individuals in at least one year; advisers that never report household clients are kept as a benchmark

[5] These records also underlie a large literature on financial misconduct, which uses the incidence of disclosures to study the prevalence, punishment, and contagion of adviser misbehavior (Dimmock, Gerken, and Graham, 2018; Egan, Matvos, and Seru, 2019).

group. We also retain one observation per firm per fiscal year, based on the latest filing.

Table I reports summary statistics for both our sample and institutional advisers (i.e. advisers that never report household clients). The median adviser in our sample is predominantly household-facing, drawing 94.3% of its client assets from households. In terms of size, sample advisers manage a median of $321 million in regulatory assets, while institutional advisers are considerably larger at $986 million; both distributions are heavily right-skewed, with means more than an order of magnitude above the medians. In terms of compensation, almost every adviser charges AUM-based fees, 98% in the sample and 94% among institutional advisers, and within the sample fixed fees (61%) and hourly fees (44%) are also quite common. For institutional advisers this is not the case: fixed and hourly fees are rare, at 13% and 2%, and performance-based fees are instead more common, at 79% against 16% in the sample. In this sense, the compensation structure reflects the clientele: advising households is primarily a relationship service, while managing institutional money is performance-compensated asset management.

2 The Demand Side: Role of AI for Households

This section examines the role of AI for households who demand financial advice. First, we show that the advice at the center of customer disputes overwhelmingly pertains to the technical aspects of financial advice such as security selection, rather than relational services such as behavioral coaching. Second, we show that AI can successfully evaluate the quality of human advice. Finally, we show that AI can also serve as a credible outside option.

2.1 The Composition of Financial Advice

Financial advice bundles technical and relational services. Not only do advisers select securities and construct portfolios, but they also coordinate and sustain the relationships that make households willing to delegate (Gennaioli, Shleifer, and Vishny, 2015). Importantly, AI does not reach these two components in the same way. While it excels where advice can be checked against models, it offers no substitute for the relationship itself (Autor and Thompson, 2025; Garicano, Li, and Wu, 2026). The scope for AI to serve households therefore depends on where their problems with advice are concentrated, which we begin by measuring.

To measure this composition, we use customer disputes from BrokerCheck and IAPD, the public disclosure registries maintained by FINRA and the SEC for brokers and investment ad-

visers. Each disclosure records a customer's formal allegation against a financial professional, including the product involved, the alleged damages, the narrative of the complaint, and its eventual resolution. The registries collectively cover the near-universe of registered advisers in the United States. We read the underlying narratives to identify which advisory service each dispute concerns. To do this, we download the disclosure reports as PDFs, convert them to text, and extract each customer-dispute event together with its allegation narrative. We focus for now on disputes reported since 2015, roughly 24,000 events in total.

With the extracted corpus, we classify each dispute narrative into a taxonomy of advisory services using GPT-5.5. We choose the categories based on the advisory activities that advisers report in Item 5.G of Form ADV, organized by function. We then augment this list with the relational and behavioral services that the form leaves out, resulting in five broad and twenty-six narrow categories. Finally, we sort the narrow categories into technical, mixed, and relational services: a technical service can be checked against an external benchmark, such as prices, fee schedules, or suitability standards, while a relational service is relationship-specific.

Table II reports the results. At the broad category level, disputes are predominantly concentrated in investment and portfolio management, which accounts for 71.9% of cases, followed by financial planning at 26.7%. Relationship and behavioral management, on the other hand, accounts for less than 1%. The narrow categories provide more detail. Security selection alone accounts for 61.0% of all disputes, followed by risk and insurance planning at 23.9%, while relational categories such as behavioral coaching and dedicated access are nearly absent. Overall, categories classified as technical account for 72.6% of disputes, mixed categories for 26.4%, and relational categories for only 1.0%.

These shares describe disputed advice, not the entire sample of advice households receive. The distinction is nonetheless useful for our purpose. These disputes isolate the episodes in which the stakes were high enough for households to seek formal recourse, which are precisely the episodes in which an outside check on advice would have been most valuable. The fact that these episodes are overwhelmingly technical indicates where the potential for AI is largest.

2.2 AI as an Auditor

Having documented that disputed advice is concentrated on the technical margin, we next ask whether AI can serve as an auditor of that advice. That is, we ask whether, given only what a client could observe, AI recognizes poor advice as poor. Our design, summarized in Figure 2, proceeds in three steps: (i) reconstruction, (ii) masking, and (iii) evaluation. The design parallels the audit study of Mullainathan, Noeth, and Schoar (2012), in which trained

auditors posed as customers and evaluated the advice they received, with AI in the role of the auditor.

Step 1: Reconstruction. A first LLM decomposes each dispute record into two objects. The first object is the client's *situation*: a first-person description of the client's circumstances at the moment advice was sought (age, employment, objectives, holdings, and constraints), restricted to facts knowable at that time. The second object is the adviser's *recommendation*: the advice at issue, with its substance preserved and ex-post evaluative labels removed. Claims the adviser made when selling the recommendation, such as describing a product as safe or guaranteeing a return, are preserved with attribution, since they are part of the advice. However, characterizations added after the fact by either party, such as "unsuitable" or "agreed upon," are removed.

Step 2: Masking. We then remove all information that could reveal the dispute or its outcome. Resolution language (settlements, awards, dismissals) is stripped, dollar amounts are masked, and product and issuer names are generalized to informative characteristics. As a result, what remains is the information set of a household evaluating advice in real time: the client's circumstances and the recommendation itself, with no trace of how events unfolded.

Step 3: Blind evaluation. A judge model then receives the situation and the recommendation and assigns the recommendation a score from 1 (clearly inappropriate for the client's situation) to 5 (well-suited). Importantly, the judge is drawn from a model family different from the one used for reconstruction, receives no source labels, and is instructed to evaluate substance while ignoring length and style. Each query is independent: the judge has no memory across cases and never observes that a complaint was filed, let alone how it was resolved.

If the disputed advice was genuinely poor, and if the client-observable facts contain enough information to detect it, the scores should concentrate at the bottom of the scale. Figure 3 shows that they do. The modal score is 2, assigned to 52% of cases, and another 23% of cases receive the lowest score of 1, so that 75% of the disputed advice is rated a 2 or below. At the other end, fewer than 8% of recommendations earn a 4, and essentially none earns a 5. The mean score is 2.09 on a scale anchored at 1 (clearly inappropriate) and 5 (well-suited). We also find that the lowest-rated cases display the patterns familiar from the misconduct literature (Dimmock, Gerken, and Graham, 2018; Egan, Matvos, and Seru, 2019), such as illiquid products sold to clients with documented liquidity needs. Collectively, these results show that a blinded model, given only the ex-ante information a client could observe, independently identifies the advice that culminated in formal disputes as poor.

Two features of this exercise are worth noting. First, because every case in the sample

culminated in a formal dispute, the distribution describes the left tail of advice quality, not the average advisory relationship. That is the relevant sample for our purpose: it isolates precisely the episodes in which an outside check on advice would have mattered. Second, the low scores are a property of the advice rather than of a harsh judge: the judge uses the full scale within the disputed sample, and, as the next subsection shows, the same judge assigns high scores to alternative recommendations for the very same client situations.

2.3 AI as an Outside Option

In addition to helping evaluate the advice households receive, AI can help by producing advice of its own. A household that can obtain sound advice directly from a general-purpose tool holds an outside option to the advisory relationship, and the strength of that outside option shapes what advisers can charge.

A growing literature suggests that the advice AI produces is good enough to make this outside option credible, and retail investors already use general-purpose AI tools to interpret financial information at scale (Blankespoor, Croom, and Grant, 2026). First, LLMs command the underlying knowledge: they pass financial-literacy tests and professional certification exams (Niszczota and Abbas, 2023; Callanan et al., 2024). Second, the advice they generate is sound: their recommendations are consistent with textbook principles of diversification and risk-matching (Fieberg, Hornuf, and Streich, 2023; Lo and Ross, 2024), and households simulated to follow LLM advice over the life cycle move toward the choices of a calibrated life-cycle model (Choukhmane et al., 2026). Finally, the technology performs at a professional level: on technical tasks such as earnings forecasting and stock analysis, LLMs match or outperform finance professionals (Lopez-Lira and Tang, 2023; Kim, Muhn, and Nikolaev, 2024; Cao et al., 2024). These capabilities are also consistent with past findings that algorithmic advice can improve outcomes for under-diversified investors and broaden access to wealth management (D'Acunto, Prabhala, and Rossi, 2019; D'Acunto and Rossi, 2023; Reher and Sokolinski, 2024).

To test whether AI can serve as an outside option in our setting, we extend the pipeline of the previous subsection with a production branch (Figure 2). For each dispute, an LLM receives only the reconstructed client situation and produces its own recommendation. Crucially, the model never sees the adviser's recommendation, nor any indication that a dispute exists. The output is style-normalized so that authorship cannot be inferred from tone or format. The blind judge then scores the two recommendations, human and AI, in randomized order without source labels. Because the judge is drawn from a model family different from the producer's, the design mitigates concerns about models favoring their own output. The comparison is

within-case: both recommendations respond to the same client situation and are scored by the same judge.

If AI advice is a credible alternative, its scores should sit well above those of the disputed human advice from the same situations. Figure 4 shows that they do, on axes identical to Figure 3. The distribution is concentrated at the top of the scale: 49% of AI recommendations score a 4 and another 42% score a 5, so that 91% are rated well-suited or nearly so, while only 2% score a 2 or below. The mean score is 4.31, against 2.09 for the human advice given the identical situations. In a direct comparison, the judge prefers the AI recommendation in 95% of cases. The result is robust to replacing the entire generation-side model family. Collectively, the evidence indicates that a general-purpose AI, given only what the client could describe about her own circumstances, produces advice that a blinded evaluator rates far above the advice these households actually received.

3 The Supply Side: AI Adoption by Advisers

In this section, we examine the supply side by studying how financial advisers themselves use AI. We focus on two questions: (i) where in the production of advice the technology enters, and (ii) how far adoption has spread. Together with the earlier results on the demand side, the answers to both questions discipline how we model the impact of AI in this market.

3.1 Measuring AI Adoption

We measure adoption of AI from Form ADV Part 2 brochures. Every registered adviser must file a narrative brochure describing its advisory services, which the adviser must update annually and amend promptly after material changes. The brochure is written to be read by prospective clients and thus describes the services a firm sells, in the firm's own words. A firm that incorporates AI into those services therefore has both the occasion and the regulatory incentive to disclose the change. To measure this adoption, we proceed in two steps.

In the first step, we flag every brochure that contains an AI-related term. The terms fall into three types: (i) core AI concepts and their acronyms, such as artificial intelligence, AI, machine learning, neural networks, and large language models; (ii) names of general-purpose AI products, such as ChatGPT, OpenAI, and Anthropic; and (iii) names of AI tools built specifically for advisers, such as meeting assistants and compliance software.

In the second step, we determine whether each flagged firm affirmatively incorporates AI

into its advisory business, as opposed to merely discussing AI as a risk factor, as an investment theme, or as a technology embedded in a vendor's product. We first classify each flagged brochure with an LLM, which supports every substantive classification with a verbatim excerpt from the filing. We then manually verify the classifications against these excerpts and the underlying filings. For the brochures that describe use, we record where in the business the technology enters, from portfolio construction and research to client-facing and back-office functions.

Throughout, we say an adviser *adopts* AI in a given year if it files a brochure that affirmatively describes using AI in its advisory services. This definition captures use that is material enough to enter the firm's own description of its business. It does not capture informal use that never reaches the brochure, such as an adviser summarizing meetings or drafting notes with an LLM. In the model, we also distinguish these two margins when we introduce the technology's effect on production and servicing costs.

3.2 The Extent and Nature of Adoption

We first summarize the extent of adoption. Figure 5 plots the share of advisers in our sample that adopt AI in each year. Adoption remains below 0.5% through 2023, doubles to 0.9% in 2024, and reaches 3.3% in 2025. In counts, the number of adopting advisers rises from 44 in 2021 to 350 in 2025. Adoption is therefore real and accelerating, but it remains early.

We next turn to where the technology enters the advisory business. Panel (a) of Table III first reports which business functions of the firm the use brochures place AI in. The investment process appears in 72.0% of use brochures and research in 50.1%. Client-facing functions appear in 38.1%, primarily focused on chat interfaces that deliver investment guidance or meeting summaries rather than AI producing the relationship itself. Back-office functions appear in 35.0%, while compliance accounts for 6.0% and marketing and operations for 3.5%.

Panel (b) answers the same question in economic terms. We reclassify the same brochures by the component of advice the described use produces, which is the distinction at the center of the paper. We find that technical production appears in 80.5% of use brochures, followed by operations and compliance at 35.2%. The client-facing uses of Panel (a) resolve into technical delivery (15.4%) and communication mechanics (12.8%): AI at the client interface delivers technical content or handles the logistics of interaction. Relational production appears in 1.3% of use brochures. In short, the supply side mirrors the demand side: disputes concentrate on the technical component and so does adoption by the advisers.

While AI adoption among registered advisers is real and increasing from a small base over our sample period, the speed of adoption is slow. For example, by late 2025, roughly 18 percent of U.S. businesses reported using AI in some business function, compared with 3.3 percent of advisers in our sample (U.S. Census Bureau, 2026). Although disclosure in a regulatory filing is a stricter threshold than a survey response, the gap indicates that adviser adoption lags the economy at large. Adoption is also slow relative to the technology's own diffusion on the demand side, where the same models reached one hundred million consumers within two months of release (OpenAI, 2026). The slow pace is perhaps not surprising because adopting AI requires complementary changes to how advice is produced, and its returns are uncertain, so waiting has value (David, 1990; Hall and Khan, 2003; Brynjolfsson, Rock, and Syverson, 2021). In addition, advisers owe their clients a fiduciary duty, so a firm that incorporates AI into its recommendations remains fully liable for them. This lag is nonetheless important to note, as it indicates that the market-level consequences of the technology are only beginning to be realized.

3.3 The Cross-Section of Adoption

Having documented that advisers use AI primarily in investment and research, we next ask which advisers are most likely to adopt it. Because the uses we observe are concentrated in technical production, the scope for adoption should scale with the proportion of technical production in a firm's business. We do not observe production mixes directly, but clientele provides a proxy: institutional clients purchase technical production largely on its own, whereas households purchase it as one component of an ongoing advisory relationship. We therefore expect adoption to decline with the household share of an adviser's client assets.

We first examine this prediction in Figure 6, which reports the share of advisers ever classified as AI users across bins of household client assets. Sample advisers are sorted into five bins of their average household share, with institutional advisers outside the sample reported as a separate comparison group. Adoption falls from 8.9% among advisers with no household clients and 8.5% in the lowest household-share bin to 2.3% in the highest, with intermediate bins between these endpoints. Overall, the figure documents a systematic clientele gradient, with lower adoption among advisers whose client assets are more concentrated in households.

One concern is that the raw gradient reflects firm characteristics correlated with clientele rather than differences in technical specialization across firms. In particular, institution-facing advisers are generally larger, and larger firms may be more likely to adopt technologies of all kinds. To address this concern, we estimate the following cross-sectional specification, which

progressively controls for scale, organizational capacity, and the adviser's business model:

$$\text{AdoptedAI}_j = \beta\, \text{HouseholdShare}_j + X_j'\gamma + \lambda_{s(j)} + \varepsilon_j, \tag{1}$$

where j indexes advisers in our sample, and AdoptedAI_j equals 100 if the adviser ever disclosed AI use during 2021–2025. HouseholdShare_j is the average household share of client assets for the adviser, while $\lambda_{s(j)}$ denotes headquarters-state fixed effects. The controls X_j first capture scale and capacity through log assets under management, log employees, log advisory employees, and years since the adviser's first observed filing. A second set of controls captures the firm's business model through its discretionary share of assets and the compensation arrangements reported by each firm.

Table IV reports the results. Column (1) shows that moving from a fully institutional to a fully household clientele is associated with a 6.8 percentage point lower adoption rate, against a mean of 5.0 percent. Column (2) adds state fixed effects, which leave the coefficient essentially unchanged at 6.3 and indicate that headquarters location explains little of the raw gradient. Column (3) adds the scale and capacity block, and the coefficient falls to 3.7, indicating that roughly half of the raw gradient reflects the larger size of institution-facing advisers. Column (4) adds the business-model block, which reduces the coefficient further to 2.0 while leaving it statistically significant. Among advisers with similar size, organization, and compensation structure, a fully household clientele therefore predicts a 2 percentage point lower adoption rate, which is roughly 40 percent of the mean. Overall, the results suggest that the clientele gradient is not an artifact of firm size and is consistent with lower adoption when technical advice is supplied within a broader client relationship.

In summary, the supply-side evidence displays the same asymmetry documented on the demand side, with AI affecting the technical component of advice rather than the relationship. Households can use AI to evaluate and produce technical advice, while advisers deploy it mainly in investment and research functions rather than in relational production. Taken together, these findings motivate the next section's model, in which AI changes technical production while leaving relationship-specific trust intact.

4 A Model of AI and Financial Advisers

Guided by the evidence in Sections 2 and 3, this section develops a simple model of the market for financial advice. The model adapts the "money doctors" view of Gennaioli, Shleifer, and Vishny (2015), in which advisers are paid for trust as much as for returns, by adding the

technical margin that AI reaches. The model is deliberately kept parsimonious, and in the next iteration of the paper we plan to develop a quantitative version of the model with heterogeneous advisers and households that can be calibrated to the advisory industry.

4.1 Setup

There are two agents: a household with wealth W and an adviser. The household allocates wealth between a safe asset and a risky portfolio with excess return $R = \mathbb{E}[R] + \varepsilon$, where $\varepsilon \sim \mathcal{N}(0, \sigma^2)$. The expected return depends on advice.

Adviser. The adviser produces the two components of advice through two separate inputs, η_R and η_τ:

$$\mathbb{E}[R] = \bar{R} + \eta_R, \qquad \tau_a = \eta_\tau, \qquad C(\eta_R, \eta_\tau) = \tfrac{1}{2}\kappa\eta_R^2 + \tfrac{1}{2}\theta\eta_\tau^2. \tag{2}$$

The technical input η_R raises the expected return above the baseline $\bar{R}$ at a cost governed by κ, the price of machine intelligence. The relational input η_τ builds the client's trust in the adviser, τ_a, at a cost governed by θ. We interpret the two inputs as expenditures purchased in separate markets: the technical input is bought in the market for research systems, data, and analyst capacity, while the relational input is client-facing staff and meeting time. The adviser also pays a servicing cost $\kappa\psi$ per dollar of assets managed, where ψ measures how automatable servicing is.

AI. The advent of AI is modeled as a fall in κ. Beyond the adviser's production and servicing costs, the same price governs the household's own use of the technology. The household purchases AI intensity I at cost κI. Intensity raises the return the household can earn on its own to $\bar{R} + \phi(I)$, with $\phi' > 0$, and lowers its posterior uncertainty about returns to $V(I)$, with $V' < 0$ and $V'' > 0$, i.e., acquiring additional information exhibits diminishing returns.[6] We interpret $\sigma^2 + V(I)$ as a Bayesian predictive variance: intensity buys a more precise signal about returns, so a fall in V reflects genuine information rather than manufactured confidence. These are the two capabilities documented in Section 2: $V(I)$ is AI as an auditor, and $\phi(I)$ is AI as an outside option.

Household. The household has mean-variance preferences, and trust enters by softening risk aversion:

$$\Gamma(\tau) = \frac{\gamma}{1+\tau}, \tag{3}$$

[6]The analysis requires that V be sufficiently convex, $V''(\sigma^2 + V) > 2V'^2$. A sharper signal also enlarges the household's risky position, which raises the value of further information; the diminishing returns to the signal must outweigh this feedback. See Appendix A.1.

where γ is baseline risk aversion. As in Gennaioli, Shleifer, and Vishny (2015), a trusted adviser reduces the cost of bearing risk, so effective risk aversion falls with τ. Trust takes two forms that attach to different channels: trust in the adviser, τ_a, operates when the client is advised, and trust in AI, τ_{AI}, operates when she invests on her own. When advised at fee f, the household chooses its risky share ω and AI intensity I to maximize

$$U(\omega, I) = \omega(\mathbb{E}[R] - f) - \tfrac{1}{2}\,\Gamma(\tau_a)\,\omega^2\left(\sigma^2 + V(I)\right) - \kappa I. \tag{4}$$

The optimal risky share gives demand:

$$\omega^*(f) = \frac{\mathbb{E}[R] - f}{\Gamma(\tau_a)(\sigma^2 + V(I^*))}. \tag{5}$$

The expression shows that demand rises with the net return and with trust, and falls with the fee and with uncertainty. Trust and uncertainty enter demand as scale factors, a property that plays a central role in the results below. The household sets AI intensity where a sharper signal is worth its cost, and I^* is determined before the choice between the adviser and self-investing.[7]

Outside option. If the household walks away, it keeps AI but loses the adviser: it earns $\bar{R} + \phi(I^*)$ on its own and bears risk with trust in AI rather than trust in the adviser. The household then solves the same portfolio problem as in Equation (4), with no fee, effective risk aversion $\gamma/(1 + \tau_{\text{AI}})$, and the same posterior uncertainty $\sigma^2 + V(I^*)$. Substituting the optimal risky share back into utility gives the value of self-investing:

$$\underline{U} = \frac{(1 + \tau_{\text{AI}})\left(\bar{R} + \phi(I^*)\right)^2}{2\gamma\,(\sigma^2 + V(I^*))} - \kappa I^*. \tag{6}$$

The first term is the indirect utility of a mean-variance investor, increasing in the return AI delivers and in the household's trust in AI; the second term is what the household spends on the technology. The value of self-investing determines how much the adviser can extract from the relationship: the better the household can do on its own, the less the adviser can charge.

[7] This timing keeps demand linear in the fee and rules out a hold-up loop in which a higher fee reduces the client's AI use, weakens her outside option, and permits a still higher fee.

4.2 The Advisory Fee

The adviser has market power and chooses the fee and the two inputs to maximize profit subject to the client's participation:

$$\max_{f,\eta_R,\eta_\tau} (f-\kappa\psi)\,\omega^* W - C(\eta_R,\eta_\tau) \qquad \text{subject to} \qquad U \geq \underline{U}. \tag{7}$$

She keeps the fee net of servicing on the assets she manages, pays the cost of the two inputs, and retains the client only if advised value exceeds the value of self-investing.

The optimal fee is the smaller of two candidate fees. The first is the unconstrained monopoly fee: with demand linear in the fee, as in Equation (5),

$$f^{\mathrm{m}} = \tfrac{1}{2}\left(\mathbb{E}[R] + \kappa\psi\right), \tag{8}$$

marginal cost plus a markup on the return the adviser delivers. The second is the highest fee consistent with the client's participation. We define relative trust as

$$s = \sqrt{\frac{1+\tau_{\mathrm{AI}}}{1+\tau_a}}, \tag{9}$$

which measures how much the client trusts AI relative to the adviser: when $s<1$ she trusts the adviser more, and as $s \to 1$ the machine becomes a perfect substitute. Setting advised value equal to the value of self-investing, the highest fee the client accepts is

$$f^{\mathrm{c}} = \mathbb{E}[R] - s\left(\bar{R} + \phi(I^*)\right). \tag{10}$$

The object $s(\bar{R}+\phi(I^*))$ is the client's reservation return: the return she can earn on her own with AI, adjusted by how much she trusts the machine relative to the adviser. The adviser extracts exactly the surplus above this reservation return.

Proposition 1. *The optimal fee is* $f^* = \min\{f^{\mathrm{m}}, f^{\mathrm{c}}\}$. *When the participation constraint binds, the fee decomposes as*

$$f^{\mathrm{c}} = \underbrace{\left(\eta_R - \phi(I^*)\right)}_{\text{return rent}} + \underbrace{\left(\bar{R} + \phi(I^*)\right)(1-s)}_{\text{trust premium}}, \tag{11}$$

the sum of a return rent, the adviser's advantage in returns over self-directed AI, and a trust premium, the value of being more trusted than the machine.

The two components of the fee in Equation (11) respond to AI differently. The return rent compensates the adviser for delivering returns the household cannot obtain alone, and as AI

raises $\phi(I^*)$, this component is competed away. The trust premium reflects the comparison between τ_a and τ_{AI}, which the technology does not directly affect.

It is important to note that trust is valuable on both margins of the adviser's problem. Differentiating Equations (5) and (10) with respect to τ_a shows that a more trusted adviser charges a higher fee and manages more assets: trust raises both the price and the quantity of advice. Furthermore, when $s > 1$ (i.e. the client trusts the machine more than the adviser), the trust premium is negative, meaning that an adviser without a trust advantage must discount the fee to retain the client. In this case, the adviser earns only the return rent and exits once the fee no longer covers the servicing cost. For now, the analysis below focuses on the region $s < 1$, in which the adviser retains a trust advantage.

4.3 The Impact of AI

Before turning to prices, we examine how the two inputs respond to the shock, which underlies the results that follow. The adviser's first-order conditions show that both inputs rise as machine intelligence becomes cheaper ($\kappa \downarrow$):

$$\frac{\partial \eta_R^*}{\partial \kappa} < 0, \qquad \frac{\partial \eta_\tau^*}{\partial \kappa} < 0. \tag{12}$$

The two inputs rise for different reasons. The technical input rises because AI directly lowers its cost. The relational input rises even though its cost θ does not fall: a decline in κ increases the assets the adviser manages, which raises the marginal benefit of the relational input. In this sense, AI and the relational input are complements.

Consider next the fee. The auditing term $V(I)$ does not appear in Equation (10): because $\sigma^2 + V$ scales the client's value both with the adviser and on her own, it cancels from the participation constraint. Auditing therefore has no direct effect on the fee: it can move the fee only through the inputs, not through the participation constraint itself. Differentiating the constrained fee with respect to κ yields three terms:

$$\frac{\partial f^{\text{c}}}{\partial \kappa} = \underbrace{\frac{\partial \eta_R^*}{\partial \kappa}}_{<0} \underbrace{-\left(\bar{R} + \phi(I^*)\right) \frac{\partial s}{\partial \kappa}}_{<0} \underbrace{-s\,\phi'(I^*) \frac{\partial I^*}{\partial \kappa}}_{>0}. \tag{13}$$

As κ falls, the first term raises the fee, because better return production raises the return the adviser delivers. The second term also raises the fee, because a higher relational input increases the adviser's trust advantage and with it the trust premium. Only the third term lowers the

fee, because a stronger outside option raises the client's reservation return.

Proposition 2. *The response of the fee to AI is ambiguous. As κ falls, the fee falls if and only if the outside-option force exceeds the sum of the return-production and trust forces:*

$$s\,\phi'(I^*)\left|\frac{\partial I^*}{\partial \kappa}\right| > \left|\frac{\partial \eta_R^*}{\partial \kappa}\right| + \left(\bar{R} + \phi(I^*)\right)\frac{\partial s}{\partial \kappa}. \tag{14}$$

A decline in κ thus exerts an ambiguous effect on the fee. On one hand, cheaper machine intelligence improves the adviser's returns and increases her trust advantage, both of which raise the fee she can charge; on the other hand, it strengthens the client's outside option, which lowers it. The common intuition that cheaper production must lower fees holds only when the outside-option force dominates. The same ambiguity arises in the unconstrained regime, where $\partial f^{\mathrm{m}}/\partial \kappa = \frac{1}{2}(\partial \eta_R^*/\partial \kappa + \psi)$: the servicing term reflects cost pass-through and lowers the fee, while the return-production term raises it. The ambiguity traces to the same price κ entering both sides of the market: the shock that lowers the adviser's cost of producing returns also strengthens the outside option that caps the fee.

We consider next the size of the advisory sector, measured by assets under management $\omega^* W$. Differentiating demand in Equation (5), a fall in κ moves the advised risky share through four channels: returns rise, trust rises, auditing improves, and the fee moves ambiguously. The first three channels all raise ω^*. The fee is the only channel that can lower it, and only if the fee rises. This result rests on trust and uncertainty entering demand as scale factors in Equation (5), together with both inputs and the household's information improving as machine intelligence gets cheaper: every nonprice channel then moves the amount delegated in the same direction.

Proposition 3. *If the fee weakly falls as machine intelligence gets cheaper, then assets under management rise. In other words, the advisory sector can shrink only if fees rise.*

Proposition 3 implies that lower fees and a shrinking sector cannot occur together: a lower fee is itself demand-expanding, and the three non-fee channels expand demand as well. The two AI capabilities therefore operate on different margins. Because $V(I)$ scales the advised and self-investing values symmetrically, it drops out of the participation constraint that caps the fee: auditing has no direct effect on the price of advice and operates instead on the quantity delegated. Auditing expands the amount invested with no direct effect on the fee, while the outside option disciplines the fee. AI as an auditor affects quantities; AI as an outside option affects prices.

4.4 Heterogeneity Across Advisers

We conclude the analysis by asking which advisers gain from a decline in the price of machine intelligence. This question follows a long literature on how automation creates winners and losers within occupations (Autor, Levy, and Murnane, 2003; Acemoglu and Restrepo, 2019; Autor and Thompson, 2025), and it also dominates the current discourse on AI and expert work.

Advisers can differ in θ, the cost of producing trust. Because the model features a representative adviser, the comparison across advisers is necessarily a comparative static: we compare market equilibria that differ in θ, abstracting from entry and from competition among advisers. While stylized, this exercise isolates how the same shock lands on differently positioned firms. Because optimized profit is a value function, the envelope theorem eliminates the indirect effects of θ on the adviser's choices, leaving only its direct effect on the cost of the relational input: $\partial \pi^* / \partial \theta = -\frac{1}{2}(\eta_\tau^*)^2$. Differentiating once more with respect to κ and using Equation (12) yields the following result.

Proposition 4. *The profit response to AI is ordered by the cost of trust:*

$$\frac{\partial^2 \pi^*}{\partial \kappa \, \partial \theta} = -\eta_\tau^* \frac{\partial \eta_\tau^*}{\partial \kappa} > 0. \tag{15}$$

Advisers with a higher cost of producing trust gain less in profit from a fall in κ, and lose more if profits fall.

Proposition 4, which we call *trust cover*, says that the less cheaply an adviser can produce trust, the smaller her profit gain from a decline in κ. This ordering relies on θ entering the adviser's problem only through the cost of the relational input, so that the envelope theorem applies, and on the relational input rising as machine intelligence gets cheaper, which the additive separability of the input costs ensures.

5 Conclusion

Households have long delegated the management of their wealth to financial advisers, on the understanding that the adviser supplies both technical guidance and the trust that makes households willing to act on it. AI enters this arrangement from both sides: it lowers the adviser's cost of producing technical advice, and it gives the client an independent means of checking that advice or producing a version of it on her own. As a result, the adviser still sells

the bundle, but what the household pays for within this bundle is an open question.

We provide three sets of findings on how AI reshapes this market, drawn from the U.S. registered investment advisory sector. First, on the demand side, the parts of advice that require an outside check are overwhelmingly technical. In a blinded evaluation of customer disputes, AI both identifies the contested human advice as poor and, given the same client circumstances, produces advice rated far higher. Second, on the supply side, advisers adopt AI on the same technical margin—in the investment process and in research—but slowly, with less than five percent of advisers in our sample ever disclosing use and adoption concentrated among institution-facing firms. Third, a model disciplined by these facts shows that the fee response to AI is ambiguous, that AI cannot simultaneously lower fees and shrink the advisory sector, and that advisers who produce trust cheaply are relatively insulated from the shock.

Our findings offer two broader implications. First, the incidence of the technology differs across advisers. The model predicts that a fall in the price of machine intelligence favors advisers who produce trust cheaply over those whose business rests on technical production. Adopting the technology is not the same as benefiting from it: institution-facing firms adopt because AI reaches their production, and that exposure is exactly why the ordering runs against them, while household-facing advisers have little to adopt and hold the asset the technology does not touch. Second, the mechanism extends beyond financial advice. Law and medicine share a similar structure: an expert sells checkable technical claims bundled with judgment, accountability, and a relationship. Our results suggest that AI will redistribute the returns to expertise toward experts who can be trusted, in financial advice and in expert services more broadly.

References

Acemoglu, Daron and David Autor, "Skills, Tasks and Technologies: Implications for Employment and Earnings," in "Handbook of Labor Economics," Vol. 4B, Elsevier, 2011, pp. 1043–1171.

Acemoglu, Daron and Pascual Restrepo, "Automation and New Tasks: How Technology Displaces and Reinstates Labor," *Journal of Economic Perspectives*, 2019, *33* (2), 3–30.

Autor, David and Neil Thompson, "Expertise," Technical Report, NBER Working Paper 2025.

Autor, David H., Frank Levy, and Richard J. Murnane, "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics*, 2003, *118* (4), 1279–1333.

Bartlett, Robert, Adair Morse, Richard Stanton, and Nancy Wallace, "Consumer-Lending Discrimination in the FinTech Era," *Journal of Financial Economics*, 2022, *143* (1), 30–56.

Berg, Tobias, Valentin Burg, Ana Gombović, and Manju Puri, "On the Rise of FinTechs: Credit Scoring Using Digital Footprints," *Review of Financial Studies*, 2020, *33* (7), 2845–2897.

Blankespoor, Elizabeth, Joe Croom, and Stephanie M. Grant, "Generative AI and Investor Processing of Financial Information," *Journal of Accounting and Economics*, 2026, *82* (2). Forthcoming.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson, "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies," *American Economic Journal: Macroeconomics*, 2021, *13* (1), 333–372.

Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond, "Generative AI at Work," Working Paper 31161, National Bureau of Economic Research 2023.

Buchak, Greg, Gregor Matvos, Tomasz Piskorski, and Amit Seru, "FinTech, Regulatory Arbitrage, and the Rise of Shadow Banks," *Journal of Financial Economics*, 2018, *130* (3), 453–483.

Callanan, Ethan, Amarachi Mbakwe, Antony Papadimitriou, Yulong Pei, Mathieu Sibue, Xiaodan Zhu, Zhiqiang Ma, Xiaomo Liu, and Sameena Shah, "Can GPT Models Be Financial Analysts? An Evaluation of ChatGPT and GPT-4 on Mock CFA Exams," in "Proceedings of the Joint Workshop of the 8th Financial Technology and Natural Language Processing (FinNLP) and the 1st Agent AI for Scenario Planning (AgentScen)" 2024, pp. 23–32.

Cao, Sean, Wei Jiang, Junbo Wang, and Baozhong Yang, "From Man vs. Machine to Man + Machine: The Art and AI of Stock Analyses," *Journal of Financial Economics*, 2024, *160*.

Chalmers, John and Jonathan Reuter, "Is Conflicted Investment Advice Better than No Advice?," *Journal of Financial Economics*, 2020, *138* (2), 366–387.

Chatterji, Aaron, David Holtz, Neel Rakholia, Prasanna Tambe, and Gawesha Weeratunga, "How Organizations Use AI: Evidence from ChatGPT," Technical Report, Working Paper 2026.

Choukhmane, Taha, Tim de Silva, Weidong Lin, and Matthew Akuzawa, "AI Financial Advice: Supply, Demand, and Life Cycle Implications," Technical Report, Working Paper, MIT Sloan 2026.

D'Acunto, Francesco and Alberto G. Rossi, "Robo-Advice: Transforming Households into Rational Economic Agents," *Annual Review of Financial Economics*, 2023, *15*, 543–563.

D'Acunto, Francesco, Nagpurnanand Prabhala, and Alberto G. Rossi, "The Promises and Pitfalls of Robo-Advising," *Review of Financial Studies*, 2019, *32* (5), 1983–2020.

Darby, Michael R. and Edi Karni, "Free Competition and the Optimal Amount of Fraud," *Journal of Law and Economics*, 1973, *16* (1), 67–88.

David, Paul A., "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox," *American Economic Review*, 1990, *80* (2), 355–361.

Dell'Acqua, Fabrizio, Edward McFowland, Ethan R. Mollick, Hila Lifshitz-Assaf, Katherine Kellogg, Saran Rajendran, Lisa Krayer, Francois Candelon, and Karim R. Lakhani, "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality," Working Paper, Harvard Business School 2023.

Dimmock, Stephen G., William C. Gerken, and Nathaniel P. Graham, "Is Fraud Contagious? Coworker Influence on Misconduct by Financial Advisors," *Journal of Finance*, 2018, *73* (3), 1417–1450.

Dulleck, Uwe and Rudolf Kerschbamer, "On Doctors, Mechanics, and Computer Specialists: The Economics of Credence Goods," *Journal of Economic Literature*, 2006, *44* (1), 5–42.

Egan, Mark, Gregor Matvos, and Amit Seru, "The Market for Financial Adviser Misconduct," *Journal of Political Economy*, 2019, *127* (1), 233–295.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock, "GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models," Working Paper, OpenAI 2023.

Fedyk, Anastassia, Ali Kakhbod, Peiyao Li, and Ulrike Malmendier, "AI and Perception Biases in Investments: An Experimental Study," 2025. Working Paper.

Fieberg, Christian, Lars Hornuf, and David J. Streich, "Using GPT-4 for Financial Advice," Technical Report, CESifo Working Paper 2023.

Foerster, Stephen, Juhani T. Linnainmaa, Brian T. Melzer, and Alessandro Previtero, "Retail Financial Advice: Does One Size Fit All?," *Journal of Finance*, 2017, *72* (4), 1441–1482.

Fuster, Andreas, Matthew Plosser, Philipp Schnabl, and James Vickery, "The Role of Technology in Mortgage Lending," *Review of Financial Studies*, 2019, *32* (5), 1854–1899.

Garicano, Luis, Jin Li, and Yanhui Wu, "Weak Bundle, Strong Bundle: How AI Redraws Job Boundaries," Technical Report, CEPR Discussion Paper No. 21453 2026.

Gennaioli, Nicola, Andrei Shleifer, and Robert Vishny, "Money Doctors," *Journal of Finance*, 2015, *70* (1), 91–114.

Greig, Fiona, Tarun Ramadorai, Alberto G. Rossi, Stephen P. Utkus, and Ansgar Walther, "Human Financial Advice in the Age of Automation," 2025. Working Paper.

Hall, Bronwyn H. and Beethika Khan, "Adoption of New Technology," in "New Economy Handbook," Academic Press, 2003.

Imas, Alex, "Sectoral Change, AI, and the Relational Sector," April 2026. Companion technical note to "What Will Be Scarce?".

Investment Adviser Association, "Investment Adviser Industry Snapshot 2026," Technical Report, Investment Adviser Association and COMPLY 2026.

Katz, Daniel Martin, Michael James Bommarito, Shang Gao, and Pablo Arredondo, "GPT-4 Passes the Bar Exam," *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 2024, *382* (2270), 20230254.

Kim, Alex G., David S. Kim, Maximilian Muhn, Valeri V. Nikolaev, and Eric C. So, "AI, Investment Decisions, and Inequality," 2024. Working Paper.

Kim, Alex, Maximilian Muhn, and Valeri V. Nikolaev, "Financial Statement Analysis with Large Language Models," Technical Report, Working Paper 2024.

Lo, Andrew W. and Jillian Ross, "Can ChatGPT Plan Your Retirement? Generative AI and Financial Advice," *Harvard Data Science Review*, 2024, (Special Issue 5).

Lopez-Lira, Alejandro and Yuehua Tang, "Can ChatGPT Forecast Stock Price Movements? Return Predictability and Large Language Models," Technical Report, Working Paper 2023.

Mainardi, Federico and Sangmin S. Oh, "The Price of Financial Advice," Technical Report, Working Paper 2026.

Mamaysky, Harry and Yuqi Zhang, "Investment Advisors to Individual Investors," Technical Report, Working Paper 2025.

Mo, Hongwei and Shumiao Ouyang, "(Generative) AI in Financial Economics," *Journal of Chinese Economic and Business Studies*, 2025, *23* (4), 509–587.

Mullainathan, Sendhil, Markus Noeth, and Antoinette Schoar, "The Market for Financial Advice: An Audit Study," Working Paper 17929, National Bureau of Economic Research 2012.

Niszczota, Paweł and Sami Abbas, "GPT Has Become Financially Literate: Insights from Financial Literacy Tests of GPT and a Preliminary Test of How People Use It as a Source of Advice," *Finance Research Letters*, 2023, *58*.

OpenAI, "How ChatGPT Adoption Has Expanded," https://openai.com/index/how-chatgpt-adoption-has-expanded/ 2026.

Reher, Michael and Stanislav Sokolinski, "Robo Advisors and Access to Wealth Management," *Journal of Financial Economics*, 2024, *155*.

Reuter, Jonathan and Antoinette Schoar, "Demand-Side and Supply-Side Constraints in the Market for Financial Advice," *Annual Review of Financial Economics*, 2024, *16*, 391–411.

Rossi, Alberto G. and Stephen P. Utkus, "Who Benefits from Robo-Advising? Evidence from Machine Learning," 2020. Working Paper.

U.S. Census Bureau, "Business Trends and Outlook Survey," Technical Report, U.S. Census Bureau 2026.

Wu, Shijie, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhanjan Kambadur, David Rosenberg, and Gideon Mann, "BloombergGPT: A Large Language Model for Finance," Technical Report, Working Paper 2023.

A Model Proofs

A.1 The Household Problem

Throughout the appendix we write $\Sigma \equiv \sigma^2 + V(I^*)$ for posterior uncertainty. Differentiating Equation (4) with respect to ω and setting the derivative to zero,

$$\frac{\partial U}{\partial \omega} = (\mathbb{E}[R] - f) - \Gamma(\tau_a)\,\omega\,\Sigma = 0 \qquad \Longrightarrow \qquad \omega^* = \frac{\mathbb{E}[R] - f}{\Gamma(\tau_a)\,\Sigma},$$

which is the demand curve in Equation (5). The first-order condition with respect to I equates the marginal value of a sharper signal to its cost:

$$-\tfrac{1}{2}\,\Gamma(\tau_a)\,(\omega^*)^2\,V'(I^*) = \kappa. \tag{A.1}$$

Since $V' < 0$, the left side is positive, and a fall in κ raises I^*. For this comparative static, the marginal value of intensity must be decreasing in I. Substituting ω^* into the left side of Equation (A.1) gives the marginal value

$$M(I) \;=\; -\tfrac{1}{2}\,\Gamma(\tau_a)\left(\frac{\mathbb{E}[R] - f}{\Gamma(\tau_a)\Sigma}\right)^2 V'(I) \;=\; -\frac{(\mathbb{E}[R] - f)^2\,V'(I)}{2\,\Gamma(\tau_a)\,\Sigma(I)^2}, \qquad \Sigma(I) = \sigma^2 + V(I).$$

Differentiating with respect to I, using $\Sigma'(I) = V'(I)$,

$$M'(I) = -\frac{(\mathbb{E}[R] - f)^2}{2\,\Gamma(\tau_a)} \cdot \frac{V''(I)\,\Sigma^2 - V'(I)\cdot 2\Sigma V'(I)}{\Sigma^4} = -\frac{(\mathbb{E}[R] - f)^2}{2\,\Gamma(\tau_a)\,\Sigma^3}\left[V''(I)\,\Sigma - 2\,V'(I)^2\right].$$

Hence $M'(I) < 0$ if and only if $V''(I)\left(\sigma^2 + V(I)\right) > 2\,V'(I)^2$, the sufficient-convexity condition stated in the text, which we maintain together with $V(I) \geq 0$; any solution of Equation (A.1) is then a maximum, and we assume parameters are such that an interior solution exists.

Substituting ω^* back into Equation (4) gives the value of being advised:

$$U = \frac{(\mathbb{E}[R] - f)^2}{\Gamma(\tau_a)\Sigma} - \tfrac{1}{2}\,\Gamma(\tau_a)\,\frac{(\mathbb{E}[R] - f)^2}{\Gamma(\tau_a)^2\Sigma^2}\,\Sigma - \kappa I^* = \frac{(\mathbb{E}[R] - f)^2}{2\,\Gamma(\tau_a)\,\Sigma} - \kappa I^*.$$

The value of self-investing follows from the same two steps with expected return $\bar{R} + \phi(I^*)$ in place of $\mathbb{E}[R] - f$ and trust τ_{AI} in place of τ_a; substituting $\Gamma(\tau_{\mathrm{AI}}) = \gamma/(1 + \tau_{\mathrm{AI}})$ yields Equation (6).

A.2 Proof of Proposition 1

For the unconstrained fee, the adviser maximizes $(f-\kappa\psi)\,\omega^*(f)\,W$ over f. Substituting the demand curve Equation (5),

$$(f-\kappa\psi)\,\omega^*(f)\,W = \frac{W}{\Gamma(\tau_a)\Sigma}(f-\kappa\psi)(\mathbb{E}[R]-f),$$

so the objective is proportional to $(f-\kappa\psi)(\mathbb{E}[R]-f)$. The first-order condition is

$$\frac{d}{df}\left[(f-\kappa\psi)(\mathbb{E}[R]-f)\right] = (\mathbb{E}[R]-f)-(f-\kappa\psi) = \mathbb{E}[R]-2f+\kappa\psi = 0,$$

which yields $f^{\mathrm{m}} = \frac{1}{2}(\mathbb{E}[R]+\kappa\psi)$, Equation (8).

For the constrained fee, set the value of being advised equal to the value of self-investing, using the two expressions derived in Appendix A.1:

$$\frac{(\mathbb{E}[R]-f)^2}{2\,\Gamma(\tau_a)\,\Sigma} - \kappa I^* \;=\; \frac{(1+\tau_{\mathrm{AI}})\left(\bar{R}+\phi(I^*)\right)^2}{2\gamma\Sigma} - \kappa I^*.$$

The terms κI^* cancel. Substituting $\Gamma(\tau_a) = \gamma/(1+\tau_a)$ on the left side and multiplying both sides by $2\gamma\Sigma$ leaves

$$(1+\tau_a)\,(\mathbb{E}[R]-f)^2 = (1+\tau_{\mathrm{AI}})\left(\bar{R}+\phi(I^*)\right)^2. \tag{A.2}$$

Taking the positive square root of both sides, the relevant branch since $\mathbb{E}[R]-f > 0$ and $\bar{R}+\phi(I^*) > 0$, and dividing by $\sqrt{1+\tau_a}$,

$$\mathbb{E}[R]-f = \sqrt{\frac{1+\tau_{\mathrm{AI}}}{1+\tau_a}}\left(\bar{R}+\phi(I^*)\right) = s\left(\bar{R}+\phi(I^*)\right),$$

with s as in Equation (9); solving for f gives Equation (10). Since revenue is single-peaked in f with peak at f^{m}, the adviser charges f^{m} when $f^{\mathrm{c}} \geq f^{\mathrm{m}}$ and the limit price f^{c} otherwise, so $f^* = \min\{f^{\mathrm{m}}, f^{\mathrm{c}}\}$. Finally, substituting $\mathbb{E}[R] = \bar{R}+\eta_R$ and adding and subtracting $\phi(I^*)$,

$$f^{\mathrm{c}} = \bar{R}+\eta_R - s\left(\bar{R}+\phi(I^*)\right) = (\eta_R-\phi(I^*)) + \left(\bar{R}+\phi(I^*)\right)(1-s),$$

which is the decomposition in Equation (11). ■

A.3 Input Choices

We derive the input choices in the unconstrained regime; the constrained case is analogous. Write $y \equiv \mathbb{E}[R] - \kappa\psi$ for the adviser's margin. Substituting $f^{\mathrm{m}} = \frac{1}{2}(\mathbb{E}[R] + \kappa\psi)$ into profit, using $f^{\mathrm{m}} - \kappa\psi = \mathbb{E}[R] - f^{\mathrm{m}} = \frac{1}{2}y$ and the demand curve,

$$\pi = (f^{\mathrm{m}} - \kappa\psi)\,\frac{\mathbb{E}[R] - f^{\mathrm{m}}}{\Gamma(\tau_a)\Sigma}\,W - C(\eta_R, \eta_\tau) = \frac{W y^2}{4\,\Gamma(\tau_a)\,\Sigma} - \tfrac{\kappa}{2}\eta_R^2 - \tfrac{\theta}{2}\eta_\tau^2.$$

Differentiating with respect to η_R, using $\partial y / \partial \eta_R = 1$ (since $\mathbb{E}[R] = \bar{R} + \eta_R$),

$$\frac{\partial \pi}{\partial \eta_R} = \frac{W y}{2\,\Gamma(\tau_a)\,\Sigma} - \kappa\eta_R = 0.$$

Multiplying through by $2\Gamma(\tau_a)\Sigma$ and substituting $y = \bar{R} + \eta_R - \kappa\psi$,

$$2\kappa\,\Gamma(\tau_a)\Sigma\,\eta_R = W(\bar{R} - \kappa\psi) + W\eta_R \qquad \Longrightarrow \qquad \eta_R\,[\,2\kappa\,\Gamma(\tau_a)\Sigma - W\,] = W(\bar{R} - \kappa\psi).$$

Differentiating with respect to η_τ, using $1/\Gamma(\tau_a) = (1 + \tau_a)/\gamma$ and $\tau_a = \eta_\tau$,

$$\frac{\partial \pi}{\partial \eta_\tau} = \frac{W y^2}{4\gamma\Sigma} - \theta\eta_\tau = 0.$$

The two conditions solve to

$$\eta_R^* = \frac{W(\bar{R} - \kappa\psi)}{2\kappa\,\Gamma(\tau_a)\Sigma - W}, \qquad \eta_\tau^* = \frac{W y^2}{4\gamma\theta\,\Sigma}. \tag{A.3}$$

We restrict attention to parameters for which $\bar{R} > \kappa\psi$ and $2\kappa\,\Gamma(\tau_a)\Sigma > W$, so that the technical input in Equation (A.3) is positive and finite, and to the region where profit is locally concave in the two inputs.

Both inputs fall in κ. Because each first-order condition contains the other input, we differentiate the two conditions jointly. We denote the two first-order conditions by

$$G_R \equiv \kappa\eta_R - \frac{W y}{2\,\Gamma(\eta_\tau)\,\Sigma} = 0, \qquad G_\tau \equiv \theta\eta_\tau - \frac{W y^2}{4\gamma\Sigma} = 0,$$

where $\Sigma = \Sigma(\kappa)$ falls as machine intelligence gets cheaper: $\partial\Sigma/\partial\kappa = V'(I^*)\,\partial I^*/\partial\kappa > 0$, since $V' < 0$ and $\partial I^*/\partial\kappa < 0$ by Equation (A.1). Under the model's timing, the adviser takes $I^*(\kappa)$, and hence $\Sigma(\kappa)$, as predetermined when choosing the inputs. The six partial derivatives of the

system are

$$\frac{\partial G_R}{\partial \eta_R} = \kappa - \frac{W}{2\Gamma\Sigma} > 0, \qquad \frac{\partial G_R}{\partial \eta_\tau} = -\frac{Wy}{2\gamma\Sigma} < 0, \qquad \frac{\partial G_R}{\partial \kappa} = \eta_R + \frac{W\psi}{2\Gamma\Sigma} + \frac{Wy}{2\Gamma\Sigma^2}\frac{\partial \Sigma}{\partial \kappa} > 0,$$

$$\frac{\partial G_\tau}{\partial \eta_\tau} = \theta > 0, \qquad \frac{\partial G_\tau}{\partial \eta_R} = -\frac{Wy}{2\gamma\Sigma} < 0, \qquad \frac{\partial G_\tau}{\partial \kappa} = \frac{Wy\psi}{2\gamma\Sigma} + \frac{Wy^2}{4\gamma\Sigma^2}\frac{\partial \Sigma}{\partial \kappa} > 0.$$

The first entry is positive because $\kappa > W/(2\Gamma\Sigma)$ is the maintained restriction above; the cross-partials are negative because each input raises the other's marginal benefit; and the κ-derivatives are positive because a higher κ raises input and servicing costs and, through lower I^*, raises uncertainty. Totally differentiating the system with respect to κ and solving by Cramer's rule,

$$\frac{\partial \eta_R^*}{\partial \kappa} = \frac{-\frac{\partial G_R}{\partial \kappa}\frac{\partial G_\tau}{\partial \eta_\tau} + \frac{\partial G_R}{\partial \eta_\tau}\frac{\partial G_\tau}{\partial \kappa}}{D}, \qquad \frac{\partial \eta_\tau^*}{\partial \kappa} = \frac{-\frac{\partial G_R}{\partial \eta_R}\frac{\partial G_\tau}{\partial \kappa} + \frac{\partial G_\tau}{\partial \eta_R}\frac{\partial G_R}{\partial \kappa}}{D},$$

where $D \equiv \frac{\partial G_R}{\partial \eta_R}\theta - \left(\frac{Wy}{2\gamma\Sigma}\right)^2 > 0$ is the determinant of the system, positive on the locally concave region. In both expressions the numerator is a sum of two negative terms, so

$$\frac{\partial \eta_R^*}{\partial \kappa} < 0, \qquad \frac{\partial \eta_\tau^*}{\partial \kappa} < 0,$$

which establishes Equation (12).

A.4 Proof of Proposition 2

Differentiate $f^{\mathrm{c}} = \bar{R} + \eta_R^* - s(\bar{R} + \phi(I^*))$ with respect to κ, noting that κ enters through η_R^*, through s via $\tau_a = \eta_\tau^*$, and through I^*:

$$\frac{\partial f^{\mathrm{c}}}{\partial \kappa} = \frac{\partial \eta_R^*}{\partial \kappa} - \left(\bar{R} + \phi(I^*)\right)\frac{\partial s}{\partial \tau_a}\frac{\partial \eta_\tau^*}{\partial \kappa} - s\,\phi'(I^*)\frac{\partial I^*}{\partial \kappa}, \tag{A.4}$$

which is Equation (13). We sign the three terms in turn. The first term is negative by Equation (12). For the second term, differentiating Equation (9),

$$\frac{\partial s}{\partial \tau_a} = \frac{\partial}{\partial \tau_a}\left(\frac{1+\tau_{\mathrm{AI}}}{1+\tau_a}\right)^{1/2} = -\frac{1}{2}\frac{(1+\tau_{\mathrm{AI}})^{1/2}}{(1+\tau_a)^{3/2}} = -\frac{s}{2(1+\tau_a)} < 0,$$

and $\partial \eta_\tau^*/\partial \kappa < 0$ by Equation (12), so the product is positive and the second term, with its leading minus sign, is negative. For the third term, $\phi' > 0$ and $\partial I^*/\partial \kappa < 0$ by Equation (A.1), so the term is positive. As κ falls, the first two terms push the fee up and the third pushes it down; the fee falls if and only if the third term exceeds the absolute value of the first two, which is the condition in Equation (14). ■

A.5 Proof of Proposition 3

Let $\Sigma \equiv \sigma^2 + V(I^*)$. Differentiating Equation (5) with respect to κ,

$$\frac{\partial \omega^*}{\partial \kappa} = \frac{\partial \mathbb{E}[R]/\partial \kappa - \partial f/\partial \kappa}{\Gamma(\tau_a)\,\Sigma} - \frac{(\mathbb{E}[R]-f)}{\Gamma(\tau_a)^2\,\Sigma}\frac{\partial \Gamma}{\partial \kappa} - \frac{(\mathbb{E}[R]-f)}{\Gamma(\tau_a)\,\Sigma^2}\frac{\partial V}{\partial \kappa}. \tag{A.5}$$

The three non-fee channels are signed by the results above:

$$\frac{\partial \mathbb{E}[R]}{\partial \kappa} = \frac{\partial \eta_R^*}{\partial \kappa} < 0, \qquad \frac{\partial \Gamma}{\partial \kappa} = -\frac{\gamma}{(1+\tau_a)^2}\frac{\partial \eta_\tau^*}{\partial \kappa} > 0, \qquad \frac{\partial V}{\partial \kappa} = V'(I^*)\frac{\partial I^*}{\partial \kappa} > 0.$$

As κ rises, returns fall, trust falls so effective risk aversion rises, and AI intensity falls so uncertainty rises. If in addition $\partial f/\partial \kappa \geq 0$, every term is negative, so $\partial \omega^*/\partial \kappa < 0$: assets under management rise as κ falls. The sector can shrink only if the fee term is negative and large enough to dominate, that is, only if the fee rises as κ falls. ■

A.6 Proof of Proposition 4

Write optimized profit on the locally concave branch as

$$\pi^*(\kappa,\theta) = \max_{f,\,\eta_R,\,\eta_\tau} (f - \kappa\psi)\,\omega^* W - \tfrac{\kappa}{2}\eta_R^2 - \tfrac{\theta}{2}\eta_\tau^2 \qquad \text{subject to} \qquad U \geq \underline{U}.$$

The parameter θ enters the objective only through the relational cost term, and the participation constraint does not involve θ. By the envelope theorem, the total derivative of π^* with respect to θ equals the partial derivative of the objective evaluated at the optimum—the indirect effects through the optimal choices $(f^*, \eta_R^*, \eta_\tau^*)$ vanish because each choice satisfies its first-order condition:

$$\frac{\partial \pi^*}{\partial \theta} = -\tfrac{1}{2}(\eta_\tau^*)^2.$$

Differentiating this expression with respect to κ, where now the dependence of η_τ^* on κ does matter,

$$\frac{\partial^2 \pi^*}{\partial \kappa \, \partial \theta} = -\eta_\tau^* \frac{\partial \eta_\tau^*}{\partial \kappa} > 0,$$

since $\eta_\tau^* > 0$ and $\partial \eta_\tau^* / \partial \kappa < 0$ by Equation (12). A fall in the price of machine intelligence therefore produces a more favorable profit response where relational input is cheap (low θ), which is the ordering in the proposition. ■

Table I
Summary Statistics

This table reports summary statistics at the advisory firm-year level over fiscal years 2018–2025, for registered advisers with positive client-type regulatory AUM in Form ADV Item 5.D. Panel (a) is the analysis sample: all adviser-years of firms with positive household regulatory assets in at least one sample year, where household regulatory assets are defined as the regulatory assets under management attributable to individual clients, both retail and high net worth, in Item 5.D. Panel (b) reports the institutional advisers, those that never report household clients; they are excluded from the sample and serve as the comparison group in Section 3.

Variable	N	Mean	SD	P10	P50	P90
Panel (a): Analysis Sample (Adviser-Years)						
Household share (%)	75,103	77.0	33.1	7.44	94.3	100
Assets under management ($ millions)	75,103	7,521	92,578	106	321	3,606
Compensation: AUM-based fees (0/1)	75,103	0.98	0.13	1.00	1.00	1.00
Compensation: fixed fees (0/1)	75,103	0.61	0.49	0.00	1.00	1.00
Compensation: hourly fees (0/1)	75,103	0.44	0.50	0.00	0.00	1.00
Compensation: performance fees (0/1)	75,103	0.16	0.37	0.00	0.00	1.00
Compensation: commissions (0/1)	75,103	0.03	0.18	0.00	0.00	0.00
Panel (b): Institutional Advisers (Adviser-Years)						
Household share (%)	37,710	0.00	0.00	0.00	0.00	0.00
Assets under management ($ millions)	37,710	11,437	135,120	130	986	15,719
Compensation: AUM-based fees (0/1)	37,710	0.94	0.23	1.00	1.00	1.00
Compensation: fixed fees (0/1)	37,710	0.13	0.33	0.00	0.00	1.00
Compensation: hourly fees (0/1)	37,710	0.02	0.14	0.00	0.00	0.00
Compensation: performance fees (0/1)	37,710	0.79	0.41	0.00	1.00	1.00
Compensation: commissions (0/1)	37,710	0.00	0.06	0.00	0.00	0.00

Table II
Advisory Services and the Composition of Customer Disputes (1/2)

This table reports the taxonomy of advisory services and the composition of customer disputes across these services. The taxonomy consists of five broad and twenty-six narrow categories, based on the advisory activities reported in Item 5.G of Form ADV and augmented with the relational and behavioral services that the form leaves out. Each narrow category is sorted into technical, mixed, or relational services. *Share* is the share of customer disputes from BrokerCheck and IAPD since 2015 whose primary service falls in each category, classified from the allegation narrative. Dashes denote categories with no disputes. The table continues on the next page, where the final rows aggregate dispute shares across technical, mixed, and relational categories.

Category	Description	Share (%)
Panel (a): Investment / Portfolio Management		**71.9**
Security selection	Security and fund selection	61.0
Asset allocation & portfolio design	Strategic/tactical allocation, risk targets	4.5
Account & vehicle structure	SMA, UMA, model portfolios	3.7
Portfolio monitoring & rebalancing	Monitoring, rebalancing, drift control	1.1
Specialized mandates	ESG, alternatives, institutional mandates	1.1
Tax-managed investing	Loss harvesting, asset location	0.5
Panel (b): Financial Planning & Advice		**26.7**
Risk & insurance planning	Life, disability, LTC, annuities	23.9
Tax planning	Projections, entity structure	1.2
Retirement & decumulation planning	Retirement income, withdrawal strategy	0.7
Estate, trust & wealth transfer	Estate, trust, succession	0.3
Education & major-goal funding	529 plans, goal funding	0.2
Comprehensive goals-based planning	Cash-flow and goal integration	0.2
Business & institutional planning	Owner, benefits, institutional advice	0.2
Philanthropic & charitable planning	DAFs, foundations, charitable trusts	0.0

Table II
Advisory Services and the Composition of Customer Disputes (2/2)

Category	Description	Share (%)
Panel (c): Relationship & Behavioral Management		**0.9**
Behavioral & emotional coaching	Discipline and expectation management	0.6
Communication & reporting	Meetings, statements, portals	0.3
Dedicated access / staffing	Named adviser, team, family office	0.0
Professional coordination / concierge	CPA/attorney coordination, life events	0.0
Panel (d): Advisory Intermediation / Manager Selection		**0.4**
External manager / sub-adviser selection	Due diligence, screening, recommendation	0.2
TAMP / platform intermediation	Turnkey third-party strategies	0.1
Delegated management (OCIO)	Delegation to third-party managers	0.1
Manager monitoring & replacement	Oversight, review, replacement	0.0
Panel (e): Information & Educational Products		**0.0**
Security ratings & pricing	Ratings and pricing services	0.0
Newsletters & periodicals	One-to-many market commentary	–
Market-timing services	Timing signals	–
Educational seminars & workshops	Investor education	–
Share of disputes in categories classified as technical		72.6
Share of disputes in categories classified as mixed		26.4
Share of disputes in categories classified as relational		1.0

Table III
Where AI Enters the Advisory Business

This table reports where AI enters the advisory business, based on the brochures of sample advisers that affirmatively describe using AI. Panel (a) reports the share of these brochures that place AI in each business function, and Panel (b) reports the share that place AI in each component of advice. Technical production covers analysis and recommendations, while technical delivery covers transmitting this content to clients. Communication mechanics covers scheduling, meeting summaries, and drafting; relational production covers trust, preference elicitation, and reassurance; operations and compliance covers non-advice functions. Functions and components are classified from the brochure text and are not mutually exclusive.

	Share of AI-use brochures
Panel (a): AI Use by Business Function	
Investment process	72.0%
Research	50.1%
Client-facing	38.1%
Back office	35.0%
Compliance	6.0%
Marketing and operations	3.5%
Panel (b): AI Use by Component of Advice	
Technical production	80.5%
Operations and compliance	35.2%
Technical delivery	15.4%
Communication mechanics	12.8%
Relational production	1.3%

Table IV
Determinants of AI Adoption

This table reports adviser-level regressions of an indicator equal to 100 if the adviser ever adopts AI during 2021–2025 on the household share of client assets. Column (1) reports the raw gradient; Column (2) adds headquarters-state fixed effects; Column (3) adds the scale and capacity controls; and Column (4) adds the business-model controls. Years since first filing counts years from the adviser's first Form ADV filing observed since 2015 and is capped at ten. The sample contains one observation per adviser and includes the institutional advisers. Heteroskedasticity-robust standard errors are in parentheses; ***, **, and * denote significance at the 1%, 5%, and 10% levels.

	(1)	(2)	(3)	(4)
Household AUM share	-6.84***	-6.30***	-3.72***	-2.02***
	(0.35)	(0.36)	(0.35)	(0.62)
Log AUM			0.47***	0.48***
			(0.12)	(0.12)
Log employees			1.42***	1.41***
			(0.42)	(0.42)
Log IA employees			2.01***	1.96***
			(0.47)	(0.47)
Years since first filing			-0.57***	-0.56***
			(0.06)	(0.06)
Discretionary AUM share				1.10**
				(0.53)
Performance fee				2.45***
				(0.43)
Commission				-1.73**
				(0.76)
Hourly or fixed fee				-0.74**
				(0.37)
PM for individuals				0.53
				(0.66)
Dep. Var Mean	4.99	4.99	4.99	4.99
State FE	No	Yes	Yes	Yes
Scale/capacity variables	No	No	Yes	Yes
Business-model variables	No	No	No	Yes
Observations	20,789	20,789	20,789	20,789
R^2	0.020	0.025	0.061	0.063

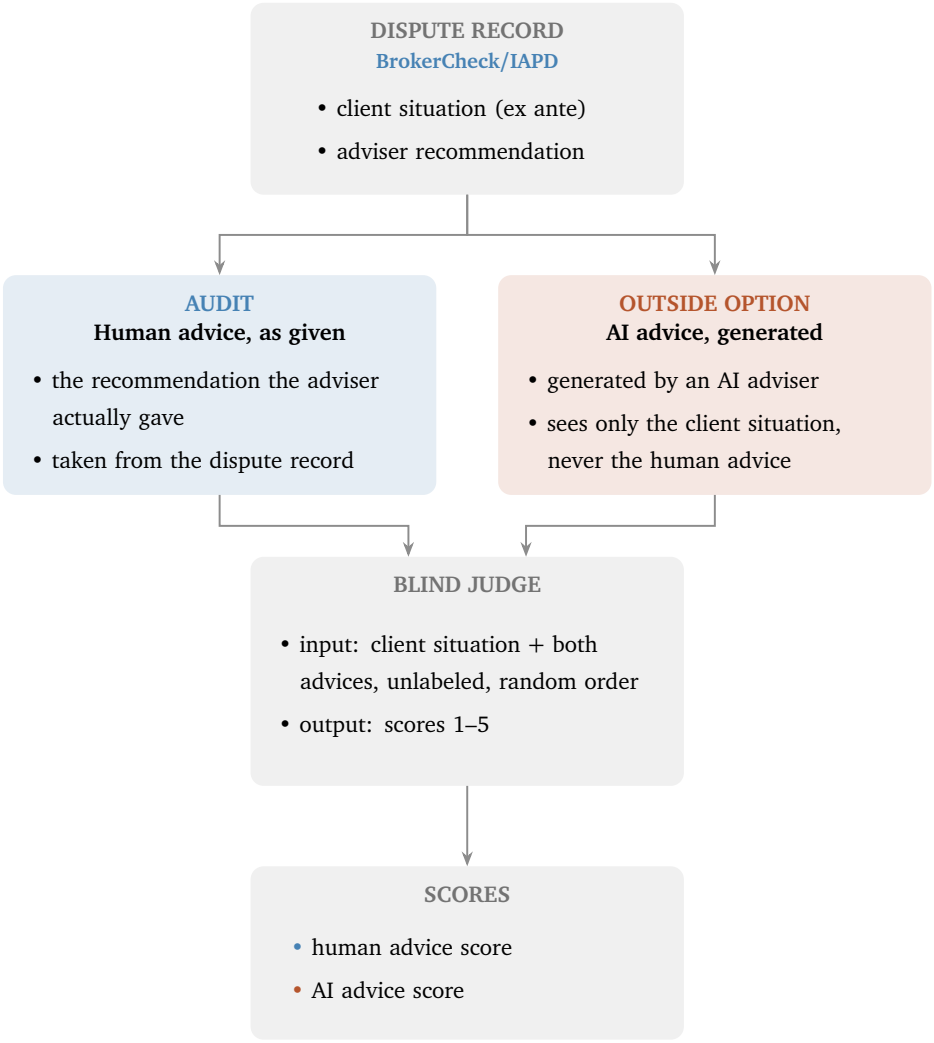


Figure 2. **Audit and Outside-Option Pipeline.** This figure summarizes the empirical design in Sections 2.2 and 2.3. Each customer dispute is decomposed into the client's situation, restricted to facts knowable when advice was sought, and the adviser's recommendation, with outcomes and identifying details masked. The human recommendation enters the audit branch as given. In the outside-option branch, an AI adviser observes only the client situation and generates its own recommendation. A blind judge observes the client situation and both recommendations in randomized order without source labels, and scores each from 1 (clearly inappropriate) to 5 (well-suited).

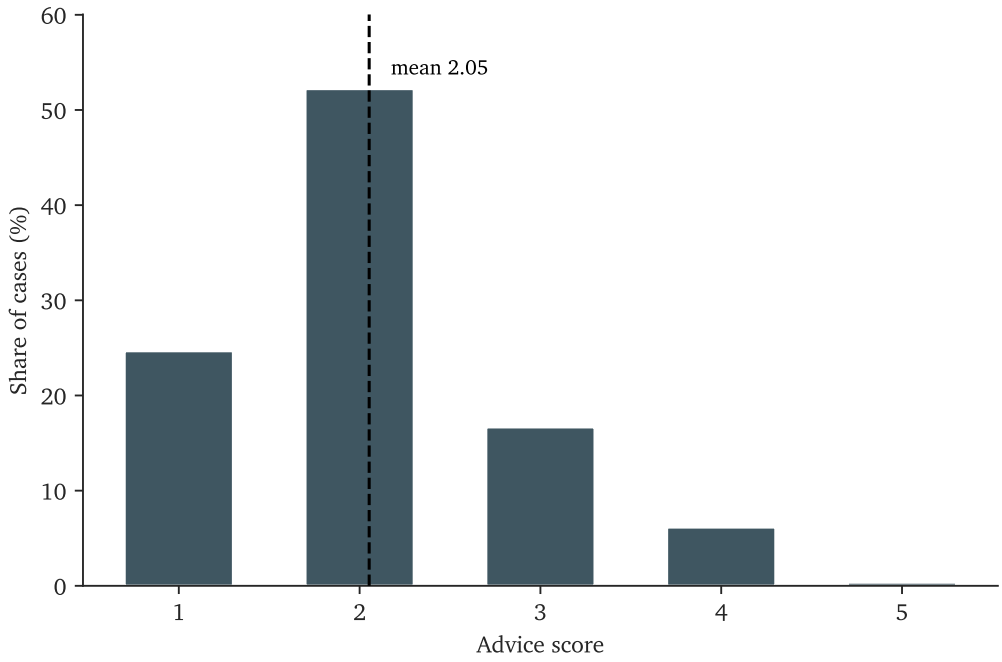


Figure 3. **Blind-Judge Scores of Disputed Human Advice.** This figure reports the distribution of blind-judge scores for the advice human advisers gave in customer-dispute cases, using the design described in Figure 2. Scores range from 1 (clearly inappropriate) to 5 (well-suited), and the dashed line marks the mean.

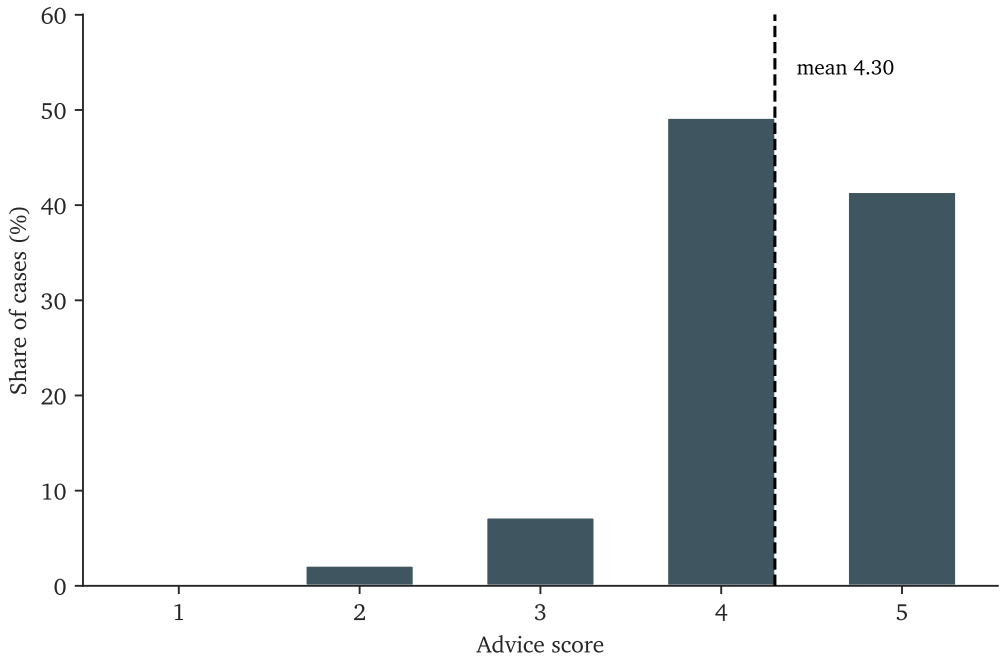


Figure 4. **Blind-Judge Scores of AI-Generated Advice.** This figure reports the distribution of blind-judge scores for AI-generated recommendations on the same client situations underlying Figure 3, on identical axes. For each case, the AI adviser observes only the client situation and produces its own recommendation (Figure 2).

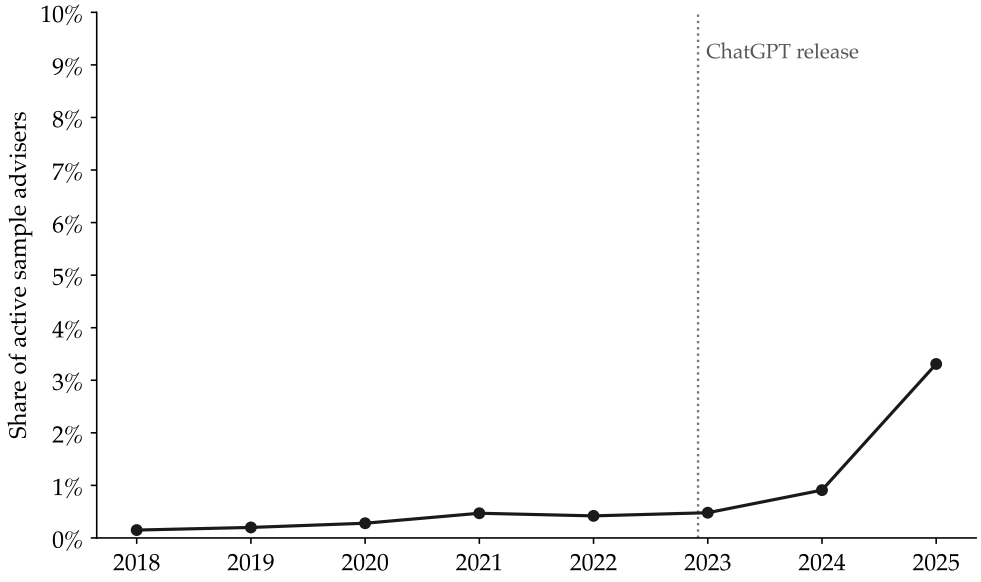


Figure 5. **The Diffusion of AI Adoption Among Advisers.** This figure plots the share of sample advisers adopting AI by fiscal year: the number of active sample advisers filing a brochure that affirmatively describes AI use, divided by the number of active sample advisers in that year (Section 1.2). The dotted line marks the release of ChatGPT in November 2022.

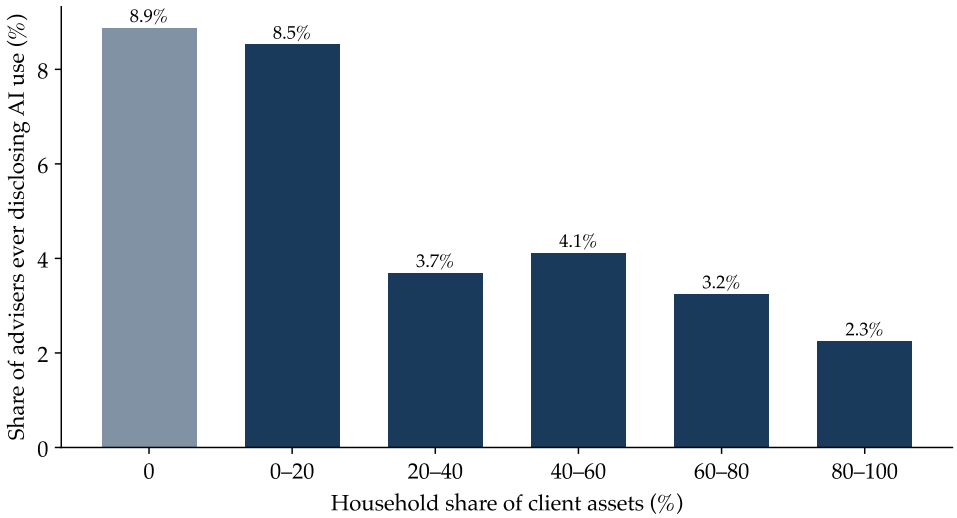


Figure 6. **AI Adoption by Client Base.** This figure plots the share of advisers ever classified as using AI in their advisory business, by the household share of client assets. The leftmost bar reports the institutional advisers outside the sample, those with no household clients; the remaining bars sort sample advisers into five bins of the average household share over the sample years. A firm counts as an AI user if any of its brochures, in any observed year, affirmatively describes using AI.