

# Artificial Intelligence and Customer Acquisition in Retail Financial Services: Experimental Evidence from Insurance Distribution

Xing Liu\*

April 2024

## Abstract

Jobs of intermediary labor force in financial services industry often involve both customer acquisition and screening. Advances in technology, however, may create information disparities across different tasks. I study how artificial intelligence (AI) affects attention allocation and information production in human-intermediated markets by analyzing a large-scale randomized experiment conducted by a top insurance agency in China. In the experiment, the firm provided treated agents with an AI-generated estimation of consumer demand for insurance, based on consumer digital footprints on the advertisements on social media; these footprints were available to all agents prior to the experiment. I show that AI demand prediction shifts agents' attention to converting high-intent consumers, improving agents' sales by 14%. As an unintended consequence, AI-generated demand information reduces agents' own information acquisition and increases adverse selection, consistent with attention models and a crowding out of risk information. Moreover, treated agents bring in riskier consumers but do not match them to more expensive products to achieve stronger incentive compatibility. The findings suggest that a common application of AI to predict consumer demand can have side effects on human information production, market efficiency, and can exacerbate agency conflicts when intermediary agents maximize their own surplus from AI.

**Keywords:** AI, intermediary agents, labor productivity, rational inattention, information acquisition, multitasking, incentives, selection markets, insurance, InsurTech, advertising, digital footprints, adverse selection, side effects

---

\*Tsinghua University PBC School of Finance. Email: [liuxing@pbcfsf.tsinghua.edu.cn](mailto:liuxing@pbcfsf.tsinghua.edu.cn). The paper was previously circulated under the title "Artificial Intelligence and Information Production in Selection Markets: Experimental Evidence from Insurance Intermediation". I am deeply grateful to Jan Bena, Will Gornall, and Sabrina Howell. I thank Alexandru Barbu, Tobias Berg, Martin Boyer, Kim Fe Cramer, Francesco D'Acunto, Joan Farre-Mensa, Lorenzo Garlappi, Shan Ge, David Goldreich, Elisabeth Honka, Zhongchen Hu, Yan Ji, Lawrence Jin, Nicola Limodio, Matt Lowe, Paolina Medina, Sendhil Mullainathan, Paul Povel, Yi Qian, Alon Rubinstein, Orly Sade, Justin Sydnor, Catherine Tucker, Jacqueline Volkman-Wise, Yongxiang Wang, Hayong Yun, Yu Zhang, Qifei Zhu, and conference and seminar participants at NBER Summer Institute, CICF, Annual Meeting of the American Risk and Insurance Association, FSU Risk Management Research Symposium, UGA PhD Symposium on Risk Management and Insurance, EIOPA-ECB Research Workshop on Insurance Companies and Pension Funds, Annual Meeting of the Society for Experimental Finance, Five Star Workshop, Asia Meeting of the Econometric Society, Australasian Finance and Banking Conference, Global AI Finance Research Conference, WEFIDEV Workshop in Finance and Development, Nanyang Technological University, Hong Kong University of Science and Technology, Chinese University of Hong Kong, Chinese University of Hong Kong (Shenzhen), Tsinghua University PBC School of Finance, Shanghai Advanced Institute of Finance, and Fudan University. I also thank Chenyang (Jason) Wei for his tremendous support when I visited Tsinghua PBCSF China Insurance and Pension Research Center. Finally, I thank the insurance agency for providing access to the data. I acknowledge the financial support from UBC Go Global Self-Directed Research Award. The views expressed in this paper are those of the author and do not necessarily reflect the views of the data provider. All remaining errors are my own.

# 1 Introduction

With the rising automation of information processing tasks, efficient attention allocation in decision-making at work becomes increasingly important for understanding labor productivity dynamics (Autor, Levy, and Murnane, 2003; Deming, 2021; Caplin et al., 2023). Incentive contracts, performance measurements, and job design are among the key conventional instruments that aim to direct humans' effort choices, yet are often inefficiently designed (Holmstrom and Milgrom, 1991).<sup>1</sup> To assist human decisions, artificial intelligence (AI) aggregates decentralized information and produces data-driven insights with lower costs of information processing. However, AI often aims to assist one task or one of several job dimensions, provides only *local* decision support, thus creating information disparities across different tasks.<sup>2</sup> How might AI redirect humans' attention allocation among different tasks and affect information production in a multitasking environment? Will AI-generated information on one task *crowd in* or *crowd out* humans' own information acquisition on another? As a result, could AI mitigate or exacerbate agency frictions? This paper answers these questions by looking at the impacts of AI on intermediary labor force in financial services industry, whose jobs often involve both customer acquisition and screening.

Understanding how AI impacts the intermediary labor force in financial services industry hinges on dissecting their job design and task attributes. About 10% of the total U.S. labor force and 12% of the U.S. financial sector were employed in sales-related occupations as of 2021.<sup>3</sup> Sales agents make costly decisions about how to market and interact by acquiring and processing flows of consumer information. A key difference between selling insurance (or finance broadly) and selling traditional consumer products is that in most occasions, sales agents deal with both consumer demand and risk (or expected cost) in a multitasking environment. In insurance markets, for example, insurers want to sell more policies but also want to insure those who are less likely to claim in the future. In credit markets, banks want to originate more loans but also want to lend to those who are less likely to default.<sup>4</sup> This highlights a common feature of sales activities in selection markets (Einav, Finkelstein, and Mahoney, 2021), namely that information about both consumer demand and consumer risk are important to sellers and sales-related jobs. Heider and Inderst (2012), for example, point out that loan officers perform dual functions in banks: they prospect for loans, creating opportunities, and screen loan applicants, gathering “soft information” that informs credit approval decisions

---

<sup>1</sup>The principal-agent problem in multitasking is examined, for example, among executives (Healy, 1985), sales force (Oyer, 1998), teachers (Carrell and West, 2010), doctors (Gravelle, Sutton, and Ma, 2010), and loan officers/examiners (Agarwal and Ben-David, 2018; Dobbie et al., 2021).

<sup>2</sup>Technology segmentation is one of the reasons. For example, at the firm level, algorithms with different objectives could be developed separately by different data scientists or even by different divisions. At the industry level, firms could be specialized in different technologies and are often evaluated by investors or other funding providers under different *tracks* or *subsegments*. Another reason is that advancements in AI technologies that assist human work across various dimensions may not be synchronized.

<sup>3</sup>Data is from the Bureau of Labor Statistics. “Sales and Related Occupations” is the second largest occupation category in the U.S. (after “Office and Administrative Support Occupations”). Sales-related occupations in finance and insurance include “Securities, Commodities, and Financial Services Sales Agents” and “Insurance Sales Agents”, employing about 0.8 million workers with a total wage expenditure of about \$65 billion. See Misra (2019) for a review of the economics of selling and sales management.

<sup>4</sup>In credit markets, loan examiners/officers serve a similar role in “sales” and incentive misalignment problems are prevalent. See, for example, Hertzberg, Liberti, and Paravisini (2010), Heider and Inderst (2012), Keys et al. (2010), Keys, Seru, and Vig (2012), Berg, Puri, and Rocholl (2020), Cole, Kanz, and Klapper (2015), Qian, Strahan, and Yang (2015), Agarwal and Ben-David (2018), among others.

from direct contacting. Loan officers' prospecting effort, however, has been largely ignored in most theories of banking where a bank faces a given demand for credit and the primary role of loan officers is screening applicants and monitoring outstanding loans. Relatedly, [Berg, Fuster, and Puri \(2022\)](#) define FinTech lending as the use of technology to facilitate credit services, which encompasses two main aspects: enhanced customer-lender interactions and borrower screening and monitoring. Existing FinTech literature has predominantly focused on the impact of new technologies that enhance risk classification and screening, while largely overlooked the role of technologies in the marketing and customer acquisition stage of financial services provision.

Studies on the impact of new technologies in the financial industry have also focused primarily on households and consumers as the target of analyses, since they are the direct users of many FinTech innovations [D'Acunto and Rossi \(2023a,b\)](#). However, most consumer choices begin by engaging with a middleman (e.g., agents, brokers, and advisors) and then make a purchase or an investment decision, especially when decisions are complex or have large stakes. Although human intermediation is common in finance and insurance, we know little about intermediary agents' decision making process underlying their intermediating outcomes.<sup>5</sup>

In this paper, I investigate how AI affects intermediary agents' attention allocation and information production in human-intermediated markets by analyzing a large-scale randomized field experiment that provides insurance agents AI-generated estimates of consumer demand, based only on information available to agents. I am able to overcome two major challenges in studying how AI-generated data impacts human decision making – data and identification. First, I collaborate with a top insurance agency in China, in which insurance agents are the mobile app users on its InsurTech platform. I access proprietary and anonymous data about insurance agents' intermediating process, outcomes, and app clickstream in the mobile environment. Second, to cleanly identify the effects of AI data processing on human decision making, the same information used by AI should be available to all agents, whereas AI-generated data is only available to treated agents. Figure 1 presents agents' inputs for decision making in the treatment and control groups. In the experiment, the firm provided treated agents with an AI-based prediction of a consumer's purchase intent for insurance, based on how the consumer had responded to advertising content on the largest Chinese social network platform, WeChat – consumer digital footprints on advertisements. Agents in both the treatment and the control group had access to information about consumer response to advertising as the raw data – the same information that AI is using to generate the demand estimates. Hence, the treatment is providing AI-processed information to agents.<sup>6</sup>

How does AI-generated data impact human decision making? First, it might impact an agent's attention and time allocation. AI represents innovation in information processing by finding complex patterns in big data and predicting variables of interest ([Agrawal, Gans, and Goldfarb, 2019](#)).<sup>7</sup> Monitoring and processing the

---

<sup>5</sup>[Foerster et al. \(2017\)](#) and [Linnainmaa, Melzer, and Previtero \(2021\)](#) show that financial advisors strongly influence their clients' asset portfolio allocation and net returns. However, these findings go beyond the sales stage and focus more on the advisory role.

<sup>6</sup>This ensures that consumer response to advertising is not new to agents. If it is new, the treatment will be providing both new data and AI-processed information from new data to agents.

<sup>7</sup>See [Babina et al. \(2024\)](#) for a brief overview of the key economic properties of artificial intelligence.

flow of information is costly and requires consistent human attention. AI makes the marginal cost of information processing very low. I first examine the information treatment effects on agents' sales productivity and the underlying mechanisms. Second, does AI-generated information affect information acquisition by humans? This trade-off is important because if AI-generated information crowds out human-collected information, the resulting information loss may distort resource allocation and market efficiency. I investigate this tension between information processing and information acquisition in the insurance market, to uncover the unintended consequences of AI-assisted decision making.

The insurance market represents an ideal laboratory to study the effect of AI on human intermediation for three reasons. First, the insurance market is marked by significant information frictions between consumers and insurers. Insurance is a complex product with *aleatory* contracts that guarantee payment to the insured for contingent events. Insurance contracts require *utmost good faith*, necessitating a higher standard of disclosure compared to most other types of contracts (Richter, Schiller, and Schlesinger, 2014). Consumer demand is the starting point. Drivers of insurance demand are multiple and often act simultaneously, such as risk, risk aversion, attention, personal and social experiences, and trust in insurers, making insurance demand hard to measure (Liebenberg, Carson, and Dumm, 2012; Outreville, 2013; Hu, 2022). AI is well-suited to complex prediction problems. In the insurance sector, AI has primarily focused on prediction problems that arise in pricing and risk selection (Eling, Nuessle, and Staubli, 2022). Users of such technologies are mainly insurers, actuaries, and underwriters. The sector has also traditionally relied on humans to intermediate the sales process. Predicting demand (take-up), which is usually part of the sales process, is therefore an important domain in which to assess human intermediation frictions. An effort of AI to aggregate decentralized information about consumer demand would be economically relevant to agents in an information-intensive industry.

Second, the insurance market is a prime example of selection markets, where consumers vary not only in how much they are willing to pay for a product but also in how costly they are to the seller (Einav, Finkelstein, and Mahoney, 2021). For example, in the classic models of Akerlof (1970) and Rothschild and Stiglitz (1976), consumers' willingness to pay for insurance increases in their (privately known) risk type or expected costs. As a result, both demand volume and effective risk selection become crucial to insurers, whereas in a traditional product market, sellers do not care about *who* the buyers are. Insurance intermediaries, as matchmakers between the supply and demand sides of insurance markets, play an important role in mitigating information asymmetries in insurance markets (Cummins and Doherty, 2006). Insurance agents, in a multitasking environment, do a lot more than selling. Initial underwriting starts with an insurance agent as the *first underwriter* in the field (Rejda and McNamara, 2014). Agents collect risk information and generally have more information than the insurer about the risk characteristics of clients. Do AI demand predictions affect agents' own information acquisition around consumer risk? By examining selection, I am able to concretize human-collected information, i.e., consumer risk information, and empirically test whether and how AI-generated information impacts information acquisition by humans.

Third, the enormous size of the insurance sector makes it a very important market to study. China is the second largest insurance market in the world with \$0.7 trillion total premiums in 2021, representing 10.2% of

the global volume.<sup>8</sup> AI has a large footprint on the insurance sector and is now central to InsurTech. The two core areas where InsurTech has the potential to add value – marketing and risk assessment – match the distinguishing feature of a selection market in that demand and cost curves are tightly linked (Einav and Finkelstein, 2011).<sup>9</sup> Distribution and intermediation – the largest InsurTech segment – accounts for one-third of recent InsurTech VC deal value.<sup>10</sup>

That demand interacts with cost in a selection market raises the question of how a common application of AI to predict consumer purchase intent to improve sales (i.e., the size of a risk pool) affects selection (the quality of a risk pool). The answer is ambiguous. One hypothesis is that information of demand *crowds out* information of risk. This could be due to agents’ endogenous allocation of attention between sales-related and risk assessment-related tasks. When AI improves the return of sales by better informing agents of consumer demand, agents prioritize their earnings gain with lower effort in detecting hidden risk information, thus reducing the quality of the risk pool. An alternative hypothesis is that information of demand *crowds in* information of risk. Insurance agents exert costly effort to acquire customers as well as collect consumer risk information for underwriting. When better informed about consumer demand by AI, agents could reallocate the saved amount of time to tasks that favor humans over AI, i.e., collecting consumer risk information. Under both hypotheses, returns to AI come from a combination of sales expansion and cost reduction. I test and help reconcile these contrasting predictions.

The collaborating insurance agency in this study is one of the largest independent agencies in China and a leading InsurTech innovator operating its own mobile app for agents. The key empirical advantage lies in my capability to utilize exogenous variation in AI information provision resulting from a randomized experiment. The experiment I analyze was conducted by the agency on its mobile app from August 9 to November 16, 2021. Digital platforms enabled randomized field experiment on a large scale. Approximately 11,000 agents are randomly assigned to a treatment or a control group from the first time they enter the relevant app page during this period. Agents in the treatment group can access predicted purchase intent given by the algorithm, including the continuous purchase intent index/score and the high-, middle-, low-intent tags categorized from the raw score, while those in the control group see the old page without any algorithmic information. Importantly, agents in both the treatment and the control group had access to information about consumer response to advertising – the same information that AI is using to generate the demand estimates. Along with granular data on sales performance, agent characteristics, policy, product, claims, and app behavior, which is particularly helpful in uncovering the *process* of intermediation, I examine the intended and unintended consequences of AI data processing on the human intermediation of insurance contracts.

---

<sup>8</sup>The U.S. is the largest with \$2.7 trillion total premiums in 2021, accounting for 39.6% of global premiums. While enormous in size, China’s insurance sector is still underdeveloped – its insurance density is only \$482 (premiums per capita) and insurance penetration is 3.9% (premiums in % of GDP), compared with that of the U.S. being \$8,193 per capita and 11.7% of GDP in 2021. See this [Sigma 4/2022 World Insurance Report](#) from Swiss Re Institute.

<sup>9</sup>Bauer et al. (2021) also note that digital innovations that enrich the customer experience and enhance core operations are closely related to the fundamental economics of insurance – the demand and supply sides of insurance provision, respectively.

<sup>10</sup>Author’s calculations based on Pitchbook’s [2021 Annual InsurTech Report](#). The six segments of the InsurTech ecosystem according to PitchBook are property & casualty, health & life, commercial, distribution & intermediation, underwriting, and administration & claims. The cooperating insurance agency and its InsurTech platform in this paper fall under the distribution & intermediation segment in the InsurTech ecosystem and the finer track – AgentTech or BrokerTech.

I show that providing predicted purchase intent improves agents' sales productivity in the treatment group by 14% in terms of total premium, relative to that of the control group. How do agents achieve the productivity gains? The main mechanism is *learning* — agents learn about consumer purchase intent from AI-processed information. Holding fixed the level of AI demand predictions of consumers, treated agents can better focus on converting high-intent consumers to final sales, when they can access AI demand predictions. I find that information provision increases the sensitivity of sales to high- and middle-intent consumers while it lowers the sensitivity of sales to low-intent consumers. Agents in the control group serve as a benchmark, as the algorithm also predicts consumer purchase intent for them, although they cannot access it. Using app clickstream data, I show that agents allocate more attention to high-intent, high-score, and top-ranked consumers. Heterogeneity tests show that the learning effect is stronger when agents are less educated (below-college degree), less experienced (below-median working experience), or underperforming (below-top-quartile sales performance in the pre-treatment period), and thus benefit more from improved information processing by AI.

To ensure that the algorithm really captures valuable information in insurance demand, I conduct two additional tests. First, I examine the information treatment effects on sales composition by policy term length. I find that agents in the treatment group sell a higher share of long-term policy and premiums (policy term length over one year), suggesting that the algorithm helps the most for insurance that is harder to sell and that requires higher capacity of information processing regarding consumer response to advertisements. Second, performance gains mainly come from more sales to new clients, suggesting that the algorithm is more helpful for weak relations where information asymmetry is higher.

The information treatment occurs during the pre-sales stage, when advisory relations are formed and a significant amount of human interaction is called for; given agents' time and attention constraints, AI information provision has the potential to facilitate more active consumer engagement and accelerate the replacement of agent-consumer relations. I find that agents in the treatment group have a higher share of new visitors (i.e., those who click on an agent's posts for the first time during the experiment) and more zombie visitors (i.e., those who clicked on an agent's posts prior to the experiment but never again during the experiment). Agents realize productivity gains by forming more new relational capital, re-energizing vintage capital, and discarding used capital.

Agents wear several hats. Improved sales productivity contributes to the revenue of insurers and the agency as well as to agents' self-interests – their earnings in the form of commission income. I show that information provision improves agents' total commission income in the treatment group by 16% and average commission income per policy by 22%, relative to that of the control group. Moreover, products sold in the treatment group have a more disperse commission rate (commission/premium). This points to the possibility that agents' financial incentives may respond to the information provision of consumer purchase intent. I find that high-commission rate products are more likely to be sold to high-intent consumers in the treatment group. Graphic evidence shows that consumers in the treatment group are more active in visiting agents' posts around the sale of a high-commission rate policy than are those in the control group. This suggests that AI-generated demand

estimates encourage agents to *cherry-pick* consumers based on sales likelihood and financial reward, consistent with information-driven consumer discrimination.

But are consumers exploited by more informed agents? Directly examining whether information treatment eventually improves selection of a more suitable product for consumers would require a benchmark specifying the optimal insurance choices for a given consumer (Gurun, Matvos, and Seru, 2016). I provide suggestive evidence that accessing AI demand predictions does not realize agents' self-interests at the cost of sacrificing consumers' interests. I show that agents in the treatment group offer a more diverse set of products, suggesting expanded search efforts for more suitable products. Using policy cancellation as a measure of sales quality and product fit, I find no evidence that AI demand predictions distort agent incentives by improving sales at the cost of offering poor-fitting products.

Finally, I examine whether and how AI demand predictions affect agents' own information acquisition around consumer risk. In insurance markets, selection speaks to the extent to which information asymmetries are resolved. In the setting of this paper, information treatment effects on selection reveal whether AI-generated demand information *crowds in* or *crowds out* consumer risk information. A positive and significant correlation between risk (ex-post realization of loss, i.e., claims) and insurance coverage (i.e., insured amount) is the necessary condition for adverse selection, suggesting that high risks buy more insurance. I show that adverse selection increased in the treatment group, especially among high- and middle-intent consumers. Increased adverse selection suggests that AI-generated demand information might discourage agents' own information acquisition, resulting in a crowding out of consumer risk information. I investigate a number of non-mutually exclusive explanations for why crowding out would occur and find supporting evidence for a combination of the following mechanisms: rational inattention, weak incentives for collecting risk information, and AI's substitution of risk information acquisition.

First, rational inattention. AI demand estimates could affect agents' selective attention allocation between sales and risk assessment. When better informed about consumer demand for insurance, agents exert more effort on sales and less effort on collecting risk information. I show that the treatment effect on adverse selection is stronger when the sales profitability associated with a mobile visitor is higher. Thus, AI demand estimates not only reallocate agents' attention across their consumer base in terms of realizing sales, but also reallocate agents' effort between sales and risk assessment, per individual consumer.

Second, weak incentives for collecting risk information. Agents may fail to foresee the future costs of providing low-quality risk information in a multi-period setting – loss of relationships with insurers and potential punishment in the form of account blocking or being kicked off the platform, for example – such that AI demand estimates crowd out their effort spent on collecting risk information at present. I show that the crowding out effect is weaker among agents who sell insurance products only from a small set of insurers thus having stronger incentives to maintain their relationships with insurers.

Third, AI-generated demand information is substitutive to risk information collected by agents, thus reducing their own information acquisition. I show that the crowding out effect on human-collected information is weaker when AI demand predictions contain less clear information about consumer risk profile.

Akerlof (1970) and Rothschild and Stiglitz (1976) suggest that consumers’ willingness to pay for insurance increases in their (privately known) risk type or expected costs. It is likely that the high-intent consumers predicted by AI are also high risks, if risk is the main driver of consumers’ interests in responding to insurance-related advertisements. Does AI draw agents’ attention to the “lemon” market? I show that AI demand estimates (tag/score/rank) do not predict ex-post claim outcomes. In digital/online environments, many factors other than consumers’ own riskiness may trigger their interests in responding to insurance-related advertisements.

Does AI affect the effectiveness of pricing instruments or the implementation of incentive-compatible contracts? Although agents do not have direct pricing power, they can match riskier customers to more expensive products via product selection. Their choice set is even broader in the setting of an independent insurance agency where product menus are from multiple insurers. I show that there are no differences of risk-premium correlation between agents in the treatment and control groups. Hence, treated agents bring in riskier customers without using pricing or product selection to achieve stronger incentive compatibility (separating equilibrium).

Overall, AI-generated demand information can augment sales productivity, but can also crowd out human-collected information on risk and exacerbate agency conflicts in a multitasking environment. Adverse selection increased due to selective attention and how humans capitalize on AI-generated data under their incentive scheme. In insurance markets or selection markets more generally, AI often aims to solve problems from one side alone – either demand (i.e., marketing and sales) or cost (i.e., risk assessment). With agency frictions and incomplete contractual design, AI may facilitate *cherry-picking* for individual agents but fail to achieve *lemon-dropping* for insurers. Information loss about consumer expected cost is the unintended consequence of AI that impairs market efficiency. I highlight an important tension when evaluating returns to AI in human-intermediated markets, namely that improved information processing may deter information acquisition.

The paper proceeds as follows. Section 2 reviews the literature. Section 3 describes the institutional background on insurance distribution and intermediation, the cooperating insurance agency, and its InsurTech platform. Section 4 details the experiment. Section 5 introduces the data. Section 6 describes the empirical strategies, while Section 7 presents the empirical results. Section 8 concludes.

## 2 Related Literature

This paper contributes to the literature on how information technology affects individuals’ decision making in the financial industry in the following ways. First, the existing literature has primarily focused on households and consumers as decision makers, since they are the direct users of many FinTech innovations.<sup>11</sup> See D’Acunto and Rossi (2023a,b) for reviews of the literature. However, consumer choices are highly intermediated, warranting a deeper understanding of the middleman’s decision making process. I add to this literature by looking at how AI data processing affects individual decision makers in intermediaries, i.e.,

---

<sup>11</sup>See, for example, D’Acunto, Rossi, and Weber (2023) and Lee (2020) on household consumption; Gargano and Rossi (2022) on savings; Chak et al. (2022) on debt management; Gargano and Rossi (2018), D’Acunto, Prabhala, and Rossi (2019), and Rossi and Utkus (2021) on trading and asset allocation; Gargano, Giacoletti, and Jarnećic (2023) on house searches; and Bundorf, Polyakova, and Tai-Seale (2024) on health insurance, among others.



agents, brokers, and advisors.<sup>12</sup> Second, I uncover both the intended and unintended consequences of AI by utilizing a large-scale randomized experiment to overcome the identification challenge and identify the effects only from AI data processing (instead of mixing effects from *data* and *data processing*). Third, with the context being the insurance market, this paper is among one of the first InsurTech studies in the literature.<sup>13</sup> Big data, AI, and InsurTech are listed by [Kojien and Yogo \(2022\)](#) and [Bauer et al. \(2021\)](#) as important areas in insurance research for further exploration. Fourth, one major advantage allowing me to examine how AI could directly impact decision making in human intermediation is access to detailed and granular data on insurance agents' intermediating behaviors via a mobile app. Since insurance agents are the app users, my focus on the direct impact of AI on agents (AI-assisted choices) also differs from most existing literature that compares machine and human decisions.<sup>14</sup> Along the lines of algorithm-assisted choices, [D'Acunto et al. \(2023\)](#) estimate the cost of cultural biases in peer-to-peer lending by comparing lenders' unassisted choices with assisted choices by an automated tool. With respect to AI and matchmaking intermediation, [Buchak et al. \(2022\)](#) examine dealer-intermediation frictions in residential real estate by studying iBuyers' (machines') intermediation process through algorithmic pricing, while I uncover agent-intermediation frictions by studying how agents (humans) respond to and utilize algorithmic information.<sup>15</sup> Fifth, the FinTech literature has focused exclusively on information technologies that predict risk (e.g., credit scoring).<sup>16</sup> Recognizing the importance of human intermediation, this paper focuses on technologies aiming to mitigate information frictions faced by the sales force – a common but understudied application of AI to predict consumer demand (i.e., take-up) – and shows that AI-generated demand metrics could also have side effects on risk.<sup>17</sup> Sixth, my results speak to how artificial intelligence might affect information production and principal-agent problems in

---

<sup>12</sup>My focus on individuals in intermediaries differs from those looking at the consequences of technology adoption at the institution level, which are endogenous choices in most occasions. See, for example, [Bonelli \(2023\)](#) on data-driven VCs and [He et al. \(2023\)](#) on IT spending in banking.

<sup>13</sup>Outside of the finance and insurance context, my study joins the growing literature on the effect of algorithms on human decision making. See, for example, [Hoffman, Kahn, and Li \(2018\)](#) on job hiring, [Kleinberg et al. \(2018\)](#) on bailing decisions, and [Ludwig and Mullainathan \(2021\)](#) on criminal justice, among others.

<sup>14</sup>See, for example, [Coleman, Merkley, and Pacelli \(2022\)](#) on comparing robot analysts with human analysts; [Cao et al. \(2023\)](#) on identifying the relative advantages of human stock analysts over machines that are constructed by the authors themselves; [Lyonnet and Stern \(2022\)](#) on using machine learning to study how venture capitalists make investment decisions; [Rossi and Utkus \(2021\)](#) on comparing human asset managers with robot advisors; and [Jansen, Nguyen, and Shams \(2023\)](#) on comparing human and machine in loan underwriting decisions, among others.

<sup>15</sup>At the individual level, the matchmaking role of intermediaries has been studied a lot more in the real estate market, i.e., real estate brokers. See, for example, [Levitt and Syverson \(2008\)](#), [Hendel, Nevo, and Ortalo-Magne \(2009\)](#), [Barwick and Pathak \(2015\)](#), [Barwick, Pathak, and Wong \(2017\)](#), [Agarwal et al. \(2021\)](#), and [Gilbukh and Goldsmith-Pinkham \(2023\)](#), among others.

<sup>16</sup>For example, [Berg et al. \(2020\)](#) study the information content of consumer digital footprint for predicting consumer default.

<sup>17</sup>In the context of a private Medicare Advantage exchange, [Gruber et al. \(2021\)](#) study the impact of a ML-based decision support tool on enrollment agents. My paper differs from theirs in that 1) the ML algorithm is to estimate the overall level of willingness/intention to buy insurance at the *extensive* margin (i.e., take-up during the sales stage) while their ML-based decision support is to recommend health insurance plans at the *intensive* margin (i.e., product selection during the enrollment stage). This difference originates from the different roles of *sales* agents vs. *enrollment* agents; 2) the algorithm input is only consumer digital footprints on the advertisements on social media – information available to all the agents prior to the experiment. The treatment is AI information processing; while theirs requires customers' personal information and prescription drug history, as well as providing agents with some new information; 3) the main mechanism causing increased adverse selection in my paper is endogenous attention allocation and information acquisition between sales and risk assessment in a multitasking environment, while in theirs it is a single-tasking scenario where improved choices lead to more acute sorting – riskier consumers are matched to more generous plans; 4) my paper focuses more on agents' productivity and performance, while theirs focuses more on downstream consumer/enrollee welfare and ex-post satisfaction.

a multitasking environment (Holmstrom and Milgrom, 1991). The mechanisms through which AI affects human decision making relate to endogenous information acquisition in financial markets (Kacperczyk, Van Nieuwerburgh, and Veldkamp, 2014, 2016; Goldstein and Yang, 2017; Farboodi and Veldkamp, 2020; Dugast and Foucault, 2018, 2022).

This paper also adds to the literature on the role of sales force in retail financial markets and how sellers' agents, brokers, and advisors shape consumer decisions. See, for example, Egan (2019) on how brokers steer consumers towards high-fee inferior convertible bonds; Hastings, Hortaçsu, and Syverson (2017) on how sales force lowered the price sensitivity of demand, resulting in high equilibrium fees in Mexico's privatized pension market; Robles-Garcia (2022) on finding that mortgage brokers increase upstream competition by lowering distribution costs, and that commission rates distort brokers' advice for households; Foerster et al. (2017) and Linnainmaa, Melzer, and Previtero (2021) on showing that financial advisors largely influence their clients' asset portfolio allocation and net returns; and Bergstresser, Chalmers, and Tufano (2009), Mullainathan, Schoar, and Noeth (2012), and Christoffersen, Evans, and Musto (2013) on agency conflicts in investment advice by mutual fund brokers, among others.<sup>18</sup> In the context of insurance,<sup>19</sup> Anagol, Cole, and Sarkar (2017) conduct field experiments to assess the quality of advice provided by life insurance agents in India. They find that agents recommend unsuitable products, cater to the beliefs of uninformed consumers, and maximize their commissions. In that study, auditors *pose* as consumers and document agents' product recommendations. In the U.S. context, Egan, Matvos, and Seru (2019) show that financial misconduct is most prominent among insurance and annuity products. Egan, Ge, and Tang (2022) examine the effect of fiduciary duty on reducing conflicting interests where brokers sell high-expense products in the U.S. variable annuities market, while Tong (2022) finds that reducing brokers' kickbacks worsens the quality of broker-intermediated plans in the employer-sponsored health insurance market. Karaca-Mandic, Feldman, and Graven (2018) find that in more competitive agent/broker markets, small businesses are more inclined to provide health insurance to their employees, and at lower premiums. Barbu (2023) shows that brokers add value to financial intermediaries by providing ex-post loss sharing via consumer exploitation. I add to this literature by analyzing a large-scale randomized experiment in a *real* insurance market to study how attention allocation and endogenous information acquisition of sales force lead to more sales-driven but less risk-driven discrimination against consumers, thus impacting market efficiency. AI and predictive algorithms help researchers diagnose agency frictions and assess inefficiencies in incentive schemes. Findings from this paper also inform whether and how AI can help resolve the multidimensional principal-agent problems in financial intermediaries.

This paper adds to the burgeoning literature on the impact of AI-based technologies on labor productivity. I provide micro-level experimental evidence that addresses the identification challenges associated with endogenous AI adoption, which have been noted in the macro-level (e.g., firm- or country-level) studies, including those by Zolas et al. (2020), Acemoglu et al. (2022), Calvino and Fontanelli (2023), and Babina

---

<sup>18</sup>See Spatt (2020) for a review on conflicts of interest in asset management and advising.

<sup>19</sup>In the Financial Economics of Insurance Workshop, Ralph S. J. Koijen and Motohiro Yogo pointed out that the role of insurance advisors (agents and brokers) is understudied and is an important future direction in insurance research. See this [teaching notes](#) for more details.

et al. (2024). At the micro-level, this paper complements Brynjolfsson, Li, and Raymond (2023), who find that access to a generative AI-based conversational assistant increases the productivity of customer support agents, with large heterogeneity in effects across workers. We both documented an average of 14% increase in productivity and a stronger effect among less-experienced and low-skilled workers. Our complementarity is manifested in the type of AI technology, the nature of job tasks, the nature of AI-generated information and the information environment, multitasking or not, and the methodology.<sup>20</sup> Importantly, the setting in my paper allows me to study how AI-generated information affects humans' collection of other types of information, and how AI affects workers' attention allocation (both within their potential customer base and between marketing and risk assessment) and agency conflicts in multitasking environments with incomplete incentive design.

This paper also contributes to the empirical literature on selection in insurance markets and financial markets more broadly.<sup>21</sup> See Einav, Finkelstein, and Levin (2010) and Einav and Finkelstein (2011) for in-depth discussions and reviews. A majority of the existing literature focuses on consumer-side selection.<sup>22</sup> Only a few studies examine supply-side (i.e., insurer-side) risk selection.<sup>23</sup> To the best of my knowledge, this paper is among one of the few to empirically study the role of individual insurance agents in shaping the risk pool and market efficiency, and how it interacts with consumer-side selection from responding to advertising.<sup>24</sup> I show that attention constraints and skewed financial incentives are key agent-side frictions contributing to selection. Moreover, my study on how AI itself and AI-assisted human decisions affect selection supports Einav, Finkelstein, and Mahoney (2021)'s view that the intersection of big data, machine learning, and AI is a promising area for studying selection markets. I highlight that AI demand predictions in selection markets could generate unintended consequences on expected cost.

This paper also connects to the literature on advertising in consumer finance.<sup>25</sup> In selection markets and health insurance markets in particular, examples include Neuman et al. (1998), Mehrotra, Grier, and Dudley (2006), Cebul et al. (2011), Aizawa and Kim (2018), and Shapiro (2020). I contribute to this literature by showing whether and how AI might enable intermediaries to extract rents from advertising to consumers, and how advertising might not be an effective risk selection tool due to intermediation frictions. Previous work focuses on firm-level advertising decisions and strategies, while ignoring the fact that advertising campaigns are malleable and often delegated to the frontline sales force. Sales agents interact with customers directly and

---

<sup>20</sup>See more discussions in the conclusion section.

<sup>21</sup>The credit market is another leading example of selection markets. A partial list of studies on information asymmetries in credit markets includes DeFusco, Tang, and Yannelis (2022) on FinTech lending; Liu, Lu, and Xiong (2024) on big tech lending; Stroebel (2016) and Gupta and Hansman (2022) on mortgages; Adams, Einav, and Levin (2009) and Einav, Jenkins, and Levin (2012) on auto loans; Agarwal, Chomsisengphet, and Liu (2010) on credit cards; and Karlan and Zinman (2009) on microloans.

<sup>22</sup>See, for example, Chiappori and Salanie (2000), Finkelstein and McGarry (2006), and Fang, Keane, and Silverman (2008).

<sup>23</sup>See, for example, Bauhoff (2012), Brown et al. (2014), Newhouse et al. (2015), Carey (2017), Decarolis and Guglielmo (2017), Aizawa and Kim (2018), and Geruso and Layton (2020).

<sup>24</sup>Allen et al. (2023) study the role of brokers in selection markets. They find broker-clients in the Canadian mortgage market finance larger loans with more leverage and longer amortization than branch-clients, which are mainly due to borrowers' preference for riskier products facilitated by brokers.

<sup>25</sup>See, for example, Grullon, Kanatas, and Weston (2004), Reuter and Zitzewitz (2006), Bertrand et al. (2010), Gurun and Butler (2012), Lou (2014), Gurun, Matvos, and Seru (2016), Hastings, Hortaçsu, and Syverson (2017), and Honka, Hortaçsu, and Vitorino (2017), among others.

have significant control over advertising customization and information extraction. My findings highlight that both consumer-side selection and agent-side choice frictions matter for studying the effects of advertising-based data analytics in selection markets. When AI interacts with the incentive scheme of the sales force, there could be side effects on expected cost (risk). Consumer-side selection creates valuable information while agent-side frictions might induce information loss (i.e., crowding out of risk information). One policy implication is how regulations could prevent institutions that gain significant information advantages from big data and AI from extracting rent and exploiting consumers (D'Acunto and Rossi, 2023a), as well as reduce potential information loss arising from intermediation frictions for better risk selection. Moreover, given that predicting consumer purchase intent is a common application of AI in marketing analytics and that big data on consumer responses to advertising also exists in other product markets, this paper also offers broader implications for the application of AI in digital marketing.

This paper provides one of the first pieces of evidence on how selective attention, as one of the human intermediation frictions, matters when AI applies to selection markets. It thus contributes to the literature on attention and inattention in decision making (Stigler, 1961; Sims, 2003; Gabaix, 2014), and on endogenous attention allocation in particular. Regarding rational inattention (e.g., Bartoš et al. (2016) and Maćkowiak, Matějka, and Wiederholt (2023)), I show that agents' selective attention across their consumer base drives productivity gains, while attention allocation between sales and risk assessment contributes to the crowding out of consumer risk information.

### 3 Institutional Background

**Insurance Distribution and Intermediation.** The vast majority of insurance is distributed through an intermediary. Affiliated agents and independent agents are the two main distribution systems. Affiliated agents (exclusive or captive agents) sell products for a single insurer under the lowest level of independence. Independent agents (non-exclusive agents) deal with multiple insurers and search for products from a wider set of suppliers. The focus of this paper is on the latter independent intermediaries.<sup>26</sup> To consumers, independent agents provide information and advisory services.<sup>27</sup> To insurers, independent agents offer distribution and marketing services as well as collect underwriting/risk information (risk assessment).

In the U.S., independent agents have the largest share of the individual life insurance market (50%), followed by affiliated agents (39%), in terms of gross written premiums in 2021.<sup>28</sup> According to the China Banking and Insurance Regulatory Commission (CBIRC), as of the end of 2021, the registered insurance agent workforce in China is 6.4 million, contributing about half of life insurance gross written premiums.<sup>29</sup>

<sup>26</sup>See Cummins and Doherty (2006) and Hilliard, Regan, and Tennyson (2013) for a detailed discussion of different distribution systems.

<sup>27</sup>Consumer inertia (e.g., mistakes or suboptimal decisions) has been documented in studies on household insurance choices. See, for example, Cutler and Zeckhauser (2004), Sydnor (2010), and Koijen, Van Nieuwerburgh, and Yogo (2016), among others. That provides another reason insurance agents play an important role in consumers' insurance purchase decisions.

<sup>28</sup>The rest comes from direct response (6%) where no producers are involved, including internet sales where consumers submit online applications; and others (5%), including financial institutions, worksite, and other channels. Data is from Insurance Information Institute, accessed on August 9, 2022.

<sup>29</sup>In China, the statistics on distribution channels are calculated by whether the premium is written by individuals or

However, the average productivity per agent is far lower than that of the U.S.. Policy makers, incumbent insurers, and technology entrants in the InsurTech landscape have been aiming to improve agent productivity and professionalism.<sup>30</sup>

**What Can Insurance Agents Do for Risk Selection?** First, agents can choose whether to deal with a potentially risky customer or not, if they are concerned about potential troubles or costs in the future. *Desk rejection* is at their own discretion. Second, conditional on selling, agents can choose the amount of effort they exert on risk assessment and information collection, which are non-routine tasks. The level of customers' full disclosure largely depends on agents' guidance. Insurers make final underwriting decisions based on that information.<sup>31</sup>

**Cooperating Insurance Agency and InsurTech Platform.** The cooperating company in this paper is a large independent insurance agency in China. The agency is also a leading technology innovator in the industry with its own InsurTech platform and a mobile app for registered agents. Employing AI, big data, and blockchain technologies, it aims to empower insurance agents by providing solutions along the entire distribution process including consumer acquisition, engagement, advertising, underwriting, client management, policy service, and claims. To that end, its online sales force is one of the largest in China. The "online+offline" mode with a technology platform represents the newest form and fastest-growing force of insurance intermediaries in China, competing with other intermediaries including affiliated agents in major insurance companies and independent agents in traditional brick-and-mortar agencies. The agency covers mainly life insurance, annuities, accident insurance, health insurance, and some property & casualty insurance (e.g., homeowners insurance). Importantly, online platforms and social media generate *hard* or *hardened* information that can be codified and processed by AI. Offline human interactions produce soft information, especially information about consumer risk profile. This hybrid process serves as an ideal setting for studying how AI affects human decision making and information production.<sup>32</sup>

**Incentive Structure.** Figure 2 shows the organizational structure of the cooperating insurance agency. Registered agents can sell products from multiple insurers. The incentive structure is as follows. Regarding sales, the compensation system is commission-based; commissions are conditioned on the insurance premium and are paid at the signing of the contract. For example, an agent will earn a 1,000 RMB commission for selling an insurance policy with a 10,000 RMB premium (10% commission rate).<sup>33</sup> Regarding underwriting,

---

institutions/agencies, rather than by agents' varying degree of independence as in the U.S.

<sup>30</sup>See, for example, this [McKinsey article](#), accessed on August 9, 2022.

<sup>31</sup>See [Regan and Tennyson \(1996\)](#) and [Schiller \(2008\)](#) for a detailed discussion of the insurance marketing system and agents' role in risk assessment. Also see [Keys et al. \(2010\)](#) and [Keys, Seru, and Vig \(2012\)](#) on the screening incentives of intermediaries in the setting of prime and subprime mortgage markets.

<sup>32</sup>See [Dumm and Hoyt \(2003\)](#) for an earlier discussion on the impact of the online distribution channel and how it interacts with the traditional agent-led channel.

<sup>33</sup>See [Hoyt, Carson, and Dumm \(2005\)](#), [Cummins and Doherty \(2006\)](#), [Cummins et al. \(2006\)](#), [Cooper \(2007\)](#), [Carson, Dumm, and Hoyt \(2007\)](#), [Cheng, Elyasiani, and Lin \(2010\)](#), [Ghosh and Hilliard \(2012\)](#), and [Ma, Pope, and Xie \(2014\)](#) for studies on the use of contingent commissions in the insurance industry, and [Cupach and Carson \(2002\)](#) in particular, for examining the influence of

agents need to collect risk information from consumers and submit that to insurers for final approval. Although there is no direct remuneration for high-quality risk assessment, there are longer-term costs to agents for low-quality risk assessment. Agents may lose their relationships with insurers, or they might be kicked off the platform if the claim ratio is too high. The agency, insurers, or regulators compile a blacklist that blocks agents' app accounts or bans agents from selling. Depending on severity, agents can face legal liability for incorrect disclosures. Other costs include customer complaints during the claim process, loss of trust, reputation concerns, and heavier post-sales service for riskier customers. Figure 3 summarizes the key elements of insurance sales as a multitasking environment.

## 4 The Experiment

The experiment I analyze was conducted by the agency on its mobile app from August 9 to November 16, 2021. It was designed independently of this study by the agency.

**Algorithm Input.** WeChat is the largest Chinese social network platform with 1.27 billion monthly active users (MAUs) in 2021, ranking third worldwide after Facebook Messenger (2.91 billion MAUs) and WhatsApp (2.29 billion MAUs).<sup>34</sup> The agency's mobile app has a plethora of advertising and marketing materials for agents to use, including product information brochures, insurance plans (either customized or general), and various types of articles (e.g., insurance knowledge, news, and claim cases). Agents can either share those on WeChat Moments to the broader audience or send them to a particular WeChat contact via direct message. The left and middle snapshots in Figure 4 provide two examples. If any WeChat user clicks and reads the content and has authorized both WeChat and the app to collect such information, her visiting records will enter the database. This user then becomes a "visitor" of the agent.<sup>35</sup>

Prior to the experiment, whenever a WeChat user clicks an advertisement, the agent receives a push notification on WeChat. The agent sees which advertisement and how much time the visitor has spent on reading it. The format is as follows.

*2022-11-01 15:40:26 Visitor A has read article 1 for 30s*

*2022-10-31 10:10:45 Visitor A has read article 1 for 10m 28s*

*2022-10-31 21:20:25 Visitor B has read article 2 for 5s*

*2022-09-10 14:30:10 Visitor A has read article 3 for 2m 15s*

*2022-08-20 08:25:55 Visitor C has read article 2 for 6m 38s*

Monitoring such information is costly, and the interpretation of a visitor's insurance demand requires consistent human attention. Agents also need to filter out noise from such raw data, as sometimes a visitor

---

compensation on product recommendations made by insurance agents using hypothetical surveys.

<sup>34</sup>Data is from [Tencent 2021 Annual Report](#), [Meta 2021 Annual Report](#), and this [Business of Apps report](#), accessed on August 10, 2022.

<sup>35</sup>If a second or a third person reposts or shares the first agent's content and a user clicks the content from there, this user is also a visitor of the first agent, although she is not in the agent's WeChat contact.

may randomly open an advertisement on social media. A machine learning algorithm uses a visitor's clicking/reading behaviors as the key inputs to predict a visitor's purchase intent. The total number of features/predictors is over 1,000. The algorithm method [non-disclosure] allows for complex non-linearity and rich interactions among predictors.<sup>36</sup> The machine learning target is the overall sales probability of a visitor.<sup>37</sup> Thus, prior to the experiment, all agents had access to information about consumer response to advertising – the same information that AI is using to generate the demand estimates, which is similar to machine learning features. During the experiment, only treated agents could access the machine learning target.

**Algorithm Output.** The snapshot on the right of Figure 4 shows the app page for this experiment.<sup>38</sup> The predictive algorithm displays two pieces of information on the app: the high-, middle-, low-intent tags based on the raw predicted score; and a continuous % score computed as the 7-day-on-7-day change of the raw predicted score representing the trend of purchase intent. Visitors with a positive (negative) % score are displayed under the Active (Silent) Tab. Screen rank is not an output of the algorithm but a feature of the screen display, where visitors are displayed in descending order of the continuous % score under each tab.

**Implementation.** Agents are randomly assigned to a treatment or a control group from the first time they enter the relevant app page (called the “visitor management center”) during the experiment period – August 9 to November 16, 2021. The randomization is based on the last digit of an agent's app account ID without any other stratifications. Agents with an even (odd) number are assigned to the treatment (control) group.<sup>39</sup> When entering the page, agents in the treatment group can access predicted purchase intent given by the algorithm, including the high-, middle-, low-intent tags categorized from the raw score and the continuous % score, displayed under the Active and Silent Tabs. Agents in the control group see the old page with old tabs without any algorithmic information. Agents in the treatment group can still access all old tabs displayed next to the new tabs. Under old tabs, visitors are displayed mechanically by the recency of visiting timestamp (order of arrival).

**Sample Formation.** The experiment is open to all agents as app users. To receive the treatment, however, they have to open the app and enter the experiment page from August 9 to November 16, 2021. I focus on registered agents affiliated with the insurance agency, thus excluding those who are simply app users but also enter the relevant page.<sup>40</sup> Only registered agents can make sales using the app and earn commissions. I require their registration date to be prior to August 9, 2021, which leads to 14,270 agents. To have a meaningful window

---

<sup>36</sup>For confidentiality reasons, I cannot provide more details about feature construction or algorithm methods.

<sup>37</sup>The algorithm only predicted the overall *intent to buy* and did not predict the *product to buy*.

<sup>38</sup>This is retrieved from the sample pages in the development stage before the agency started to implement the experiment, so that the visitor profiles here are hypothetical. The picture is pixelated as per the company's request.

<sup>39</sup>The agency commonly uses this randomization method when running A/B testing for new app products/features. I confirmed from the product manager and the developer that account ID is randomly assigned when an agent registers on the app, and so is the last digit. I also confirmed that there were no other testings going on during the same experiment period.

<sup>40</sup>Everyone can download the app for free. Those who only use the app but are not registered with the agency include agents affiliated with other agencies but who also use the app to view product information, compare products in different platforms, or use certain tools on the app. Others are active app users but who do not sell insurance.

length for identifying treatment effects, I require at least four weeks (28 days) as the observational window; agents in the baseline sample thus enter the app page between August 9 and October 20, resulting in 11,125 agents in total, including 5,430 agents in the treatment group and 5,695 agents in the control group.<sup>41</sup> Figure 5 shows the timeline of the experiment.

## 5 Data and Summary Statistics

This section describes the data sources, presents summary statistics for variables of interest, and tests for balance across the treatment and control groups.

### 5.1 Data Sources

The data contains seven parts: agent characteristics; sales and commissions; policy, policyholder, and product; claims; app behaviors; consumer visiting records; and purchase intent index. The data was provided directly by the company and was collected by their backstage system. All information was anonymized and de-identified by the company to guarantee user privacy.

**Agent Characteristics.** This data contains information on all agents who registered at the agency and created a profile on the app. Main variables include agents' gender, age, education degree, city of residence, the date on which the app profile was created, the date on which the agent registered at the agency, and branch company.

**Sales and Commissions.** This data contains information on all the policies sold by each agent. For each policy, I access variables on the date of sales, premiums, and commissions paid to the agent.

**Policy, Policyholder, and Product.** This data contains information on policy and policyholder characteristics and the associated products for all sales. Main variables include insured amount, policy term length, whether and when a policy is surrendered/canceled, product, product type/lines, insurance company, gender, age, city of residence of the applicant and the insured, and relation between the applicant and the insured (self, children, parents, spouses, and others).

**Claims.** This data contains information on claims at the policy level, including the number of reported claims (total, compensated, and not compensated) and the claimed amount (reported and actually compensated). Claim data is available only for certain products in student safety insurance, accident insurance, and health insurance.<sup>42</sup>

---

<sup>41</sup>Without imposing any filters that are correlated with whether receiving the treatment or not, one possibility is that during the experiment period the number of agents entering the app page with an even-digit ID happens to be slightly smaller than that with an odd-digit ID.

<sup>42</sup>Key to the paper, this is orthogonal to treatment assignment.



**App Behaviors.** This data contains information on app usage behaviors for all agents. One observation is one click by an agent. I access variables on the click timestamp, block/page names of the app clicked or opened, device type, and device size. On the experiment page – visitor management center – I can observe which visitor displayed on screen is clicked by the agent at which time.

**Consumer Visiting Records.** This data contains information on consumer visiting records on agents' shared or posted content on WeChat. One observation is one click/read by a visitor. I access variables on the click timestamp, duration of reading in milliseconds, content title, content type, and article tags/classifiers (e.g., health, education, disease, and social insurance).

**Purchase Intent Index.** This data contains information on the algorithm output, which is essentially what an agent sees on the app every day on the experiment page – visitor management center. The data is at the agent-visitor-day level, including date, high-, middle-, low-intent tags, continuous % score, and whether the visitor is displayed under the Active or the Silent Tab. Although agents in the control group could not access this information on the app, the algorithm has predicted purchase intent for all of them.

## 5.2 Summary Statistics and Balance Checks

Table 1 presents summary statistics and balance checks for agent characteristics and pre-treatment performance across the treatment and control groups. Of the 11,125 agents in the baseline sample, 51% are female, the average age is 41 years old as of 2021 (range: 18-75), 30% have at least a college degree, and 43% are affiliated with a branch in a first-tier city (Beijing, Shanghai, Guangzhou, or Shenzhen). The average time since registering at the agency is 11 months and the average time since app registration is 27 months, as of August 2021.<sup>43</sup> 53% of agents sold at least one policy before the experiment started, with 10 policies on average. 16% of sales count and 28% of premium are from long-term policies (i.e., term length over one year). 85% of sales count and 84% of premium are from new clients (vs. former clients). The average commission rate (commission/premium) of products sold is 24%. The average claim ratio is 8%. The average cancelation ratio is 3%. Agents sold 6 unique products on average. Balance checks along those variables suggest randomization is successful.<sup>44</sup>

---

<sup>43</sup>A majority of agents on this platform have worked in the insurance industry (e.g., as affiliated agents in traditional insurance companies or as independent agents in traditional intermediaries) for a long time before joining the company. This reflects well the status quo of insurance intermediaries in China and the move towards the mature insurance intermediation system in some developed economies where independent agents lead the market share.

<sup>44</sup>Variables are well-balanced between agents in the treatment and control groups within the final analysis sample. In comparing the analysis sample to the full agent workforce of the company, however, agents in the analysis sample have higher sales performance in the pre-treatment period. Their number of policies and total premium sold prior to August 9, 2021 (starting date of the experiment) are about two times that of those who are not in the analysis sample. It is reasonable that those who did not use and open the relevant app page during the experiment period are less-active/productive sellers. Those who entered the final analysis sample are the most relevant agent sample for estimating the treatment effects. The cross-sectional result that AI helps less-experienced and less-productive agents more in better targeting high-intent consumers is reassuring, because they are more similar to the population average. To the degree that more and more rookie agents might take up the app tool for digital marketing, this paper provides lower-bound estimates on how a first attempt of AI demand prediction might affect agents' performance, information production, and adverse selection.

## 6 Empirical Approach

**Baseline Effect.** As the treatment was assigned randomly, for baseline empirical specification (Equation (1)) I regress outcome variables on an indicator for being a treated agent to study the information treatment effects on sales productivity, commission income, product provision, and app behaviors.

$$Y_a = \beta_0 + \beta_1 Treat_a + Z_a + \varepsilon_a \quad (1)$$

where  $a$  indexes an agent,  $Y_a$  is an agent’s outcome variable, and  $Treat_a$  is an indicator variable that takes the value of one if an agent is in the treatment group, and zero otherwise. The dependent variables are measured from the day agents enter the app page until November 16, 2021 (or a day earlier if they leave the platform before November 16). As the algorithm only captures consumers who have mobile visiting records (i.e., who had responded to any advertisement on WeChat) for predicting purchase intent, sales performance and commission income do not include sales to consumers who have no mobile visiting records. This applies to both the treatment and control groups.  $Z_a$  are a set of baseline control variables, including total number of policies, total premium up until the day before the agent enters the app page, and the number of days of the observational window. Figure 5 shows the time frame of variable construction. Robust standard errors are used.

**Testing Mechanisms.** To understand the information treatment effects on sales productivity, I employ a specification at the agent-visitor level by testing whether information increases the sensitivity of sales to predicted purchase intent, using Equation (2).

$$Y_{a,c} = \beta_0 + \beta_1 Treat_a \times A_{a,c} + \beta_2 A_{a,c} + \beta_3 Treat_a + Z_{a,c} + \varepsilon_{a,c} \quad (2)$$

where  $a$  indexes an agent,  $c$  indexes a visitor/customer,  $Y_{a,c}$  is an agent’s sales to a visitor during the experiment period,  $A_{a,c}$  is a visitor’s baseline purchase intent information given by the algorithm, and  $Treat_a$  is an indicator variable that takes the value of one if an agent is in the treatment group, and zero otherwise. Agents in the treatment group could access  $A_{a,c}$  while agents in the control group could not, though the algorithm has predicted for them. The model tests whether AI information provision changes the sensitivity of sales to visitors’ purchase intent by comparing the baseline sensitivity for the control group ( $\beta_2$ ) with the increase in sensitivity for the treatment group ( $\beta_1$ ). The key prediction is that information treatment makes the sensitivity higher, so that  $\beta_1 > 0$ . This approach is similar in spirit to [Dizon-Ross \(2019\)](#), who examines how providing parents with their children’s academic performance affects the parents’ educational investments.  $\beta_3$  represents the effect on sensitivity change for the omitted intent group when  $A_{a,c}$  is a category variable, and is retained only for specifications without including agent fixed effects. As in Equation (1), I focus on cross-sectional differences between agents in the treatment and control groups.<sup>45</sup> Standard errors are clustered at the agent level.

For each visitor,  $Y_{a,c}$  is measured from the day the agent sees the visitor for the first time (i.e., when the agent enters the app page and the visitor is displayed on the screen on that day) until November 16, 2021 (or

---

<sup>45</sup>Results by including agent fixed effects are presented in Appendix A. Including visitor fixed effects could provide even stronger identifications. However, the total number of observations at the agent-visitor level is close to the unique number of visitors, suggesting there are few cases where one visitor belongs to multiple agents, thus making visitor fixed effects infeasible.

a day earlier if the agent leaves the platform before November 16).  $A_{a,c}$ , including *High-intent*, *Middle-intent* (where *Low-intent* is omitted), *Score*, and *Rank* are the visitor’s baseline information seen by the agent on the first day. Tag indicators, *Score*, and *Rank* are included in separate regressions one at a time to avoid high correlations. *Score* is displayed on the app screen as a % and divided by 100 in regressions for presenting coefficients. *Rank* is capped at 100. All models include  $Z_{a,c}$  as baseline control variables at the agent-visitor level, including total number of policies, total premium up until the day before the agent entered the app page and saw the visitor, and the number of days of the observational window for each agent-visitor.

**Attention Allocation.** To test how agents allocate attention to information components on the mobile screen, I employ Equation (3) to predict the likelihood of an agent clicking a visitor on the app page.

$$Y_{a,c,t} = \beta_0 + \beta_1 A_{a,c,t} + Z_{a,c,t} + \alpha_a + \gamma_t + \varepsilon_{a,c,t} \quad (3)$$

The sample is restricted to agents in the treatment group. The unit of observation is an agent-day-visitor and includes all days an agent entered the app page during the experiment period.  $a$  indexes an agent,  $c$  indexes a visitor/customer,  $t$  indexes a day,  $Y_{a,c,t}$  is an indicator variable that takes the value of one if the agent clicked the visitor on a day, and zero otherwise.  $A_{a,c,t}$  is a visitor’s algorithmic information on a day, including *High-intent*, *Middle-intent* (where *Low-intent* is omitted), *Score*, and *Rank*. They are included in separate regressions one at a time.  $Z_{a,c,t}$  are baseline control variables at the agent-day-visitor level, including total number of policies and total premium up until the day before the agent entered the app page and saw the visitor. All models include agent ( $\alpha_a$ ), day of week, and month ( $\gamma_t$ ) fixed effects. Standard errors are clustered at the agent level.

**Selection.** To examine whether and how AI-assisted agent decisions affect selection, I employ the classical risk-coverage correlation model (Cohen and Siegelman, 2010; Eling, Jia, and Yao, 2017) in Equation (4). A positive and significant correlation between risk (ex-post realization of loss) and insurance coverage is the necessary condition of adverse selection, suggesting that high risks buy more insurance. An insignificant risk-coverage correlation suggests no adverse selection. A negative and significant risk-coverage correlation implies advantageous selection.

$$Risk_i = \beta_0 + \beta_1 Coverage_i + Z_i + \varepsilon_i \quad (4)$$

where  $i$  indexes a policy,  $Risk_i$  is measured as the logarithm of claimed amount (eventually compensated), and  $Coverage_i$  is measured as the logarithm of insurance amount. A claim dummy that takes the value of one if there is any claim for a policy  $i$ , and zero otherwise, is used for robustness check.

$Z_i$  is a vector of control variables that insurers may use for risk classification and for mitigating asymmetric information. These are essential, as any presence of adverse selection should reflect residual asymmetric information after factoring in risk classification (Dionne and Rothschild, 2014). I control for demographics of the insured – age and gender, residential province, and the relation between the applicant and the insured; since if they are not the same person, agents may only communicate with the applicant most of the time, resulting in more severe asymmetric information. Moreover, assuming price represents insurers’ best estimation of risks from an actuarial perspective, I also control for premium rate (premium/insurance amount)

(Finkelstein and McGarry, 2006). I add product type fixed effects to account for heterogeneity in different insurance markets. To examine the role of intermediation frictions in selection in the context of this paper, I further include insurer fixed effects and agent fixed effects one at a time in Equation (4) in separate regressions to see if they increase the explanatory power of claim risk. Standard errors are clustered at the agent level.

The sample for this test includes all policies sold to visitors during the experiment period where all variables used in this regression are available. As the focus is on the correlation between risk and coverage, and selection might exist in some agent-visitor cohorts but not in others, I first estimate Equation (4) separately for policies sold among agent-visitors in the control group, treatment-high-intent group, treatment-middle-intent group, and treatment-low-intent group. To compare the magnitudes of risk-coverage correlations between agent-visitor groups and gauge the information treatment effects on selection, I then estimate Equation (4) by interacting  $Coverage_i$  with  $Treat_a$  and a visitor’s baseline purchase intent information given by the algorithm,  $A_{a,c}$ , and focus on the coefficients on  $Coverage_i \times Treat_a \times A_{a,c}$ . To investigate the potential mechanisms leading to selection, I further interact the equation with variables of interest  $X$  at the agent level or at the visitor level.

## 7 Empirical Results

### 7.1 Information Treatment Effects on Sales Productivity

Table 2 reports the information treatment effects on agents’ sales productivity, using Equation (1). The dependent variables are *# Policy* – total number of policies, *Tot. Premium* – total premium, and *Avg. Premium* – average premium per policy ( $Tot. Premium / \# Policy$ ). All dependent variables are winsorized at the 1st and 99th percentiles. *Avg. Premium* is available only for 4,687 agents who sold at least one policy during the experiment period, including 2,321 agents in the treatment group and 2,366 agents in the control group. Columns (2), (4) and (6) include baseline control variables, including total number of policies, total premium up until the day before agents entered the app page, and the number of days of the observational window. Means of dependent variables in the control group are reported in the bottom row.

I show that agents in the treatment group sell on average 0.311 more policies (Column (2), 11% higher than the control group mean) and 275.733 RMB (41 USD)<sup>46</sup> more total premium (Column (4), 14% higher than the control group mean). I find no significant increase in average premium per policy (Column (6)). This suggests that AI demand estimates improve agents’ sales productivity mainly through increasing unit of sales. Omitting or including baseline controls does not change the magnitude of coefficients much, indicating the sample is well-balanced.

Given that sales outcomes for insurance follow a highly skewed non-normal distribution, Appendix Table A.1 reports the baseline effects using performance outcomes in logs (Panel A) and inverse hyperbolic sine (IHS) transformation (Burbidge, Magee, and Robb, 1988; de Mel et al., 2022). Because absolute policy counts and premiums are still the most relevant measures for agent performance, I report non-transformed effects in the main table.

---

<sup>46</sup>CNY-USD exchange rate on August 12, 2022.

Agents enter the app page on different dates during the experiment period. The timing of entry is uncorrelated with whether they will receive the treatment or not upon entry. To have a meaningful window length for identifying the treatment effects, I require at least four weeks (28 days) as the observational window in the baseline sample above. Appendix Figure A.1 reports the information treatment effects on agents' sales performance when varying the agent sample by timing of entry. The point estimates represent coefficients of *Treat* (with their 95% confidence intervals) in Equation (1) by requiring a different minimum number of weeks as the observational window; thus the longer the window, the earlier the agent entered the app page. For example, *w1* shows the effect when requiring at least one week (7 days) for observing performance. Across the three performance outcomes, I show that a longer window is necessary for detecting larger and stronger information treatment effects.<sup>47</sup> An interesting pattern is that the value addition of AI demand estimates shifts from expanding the unit of sales in the short run (Panel A) to increasing premium per unit (Panel C) in the long run. One possibility is that information of purchase intent helps agents to acquire new customers first, and over time to improve the capacity of existing client pool.

## 7.2 Mechanisms Underlying Productivity Gains

I posit that the main mechanism through which agents achieve productivity gains is *learning* — agents learn about consumer purchase intent from AI-processed information. Holding fixed the level of AI demand predictions of consumers, treated agents can better focus on converting high-intent consumers to final sales, when they can access AI demand predictions. I conduct several tests to support this *learning* channel.

**Sensitivity of Sales to Purchase Intent.** At the agent-visitor level, Table 3 reports whether AI information provision changes the sensitivity of sales to predicted purchase intent using Equation (2). Agents in the control group serve as a benchmark, as the algorithm also predicts purchase intent for them although they cannot access it. The key is to compare the baseline sensitivity for the control group ( $\beta_2$ ) with the increase in sensitivity for the treatment group ( $\beta_1$ ). The dependent variables are *Any Policy* – a dummy variable indicating whether the agent sold any policy to a visitor, *# Policy* and *Tot. Premium*. I estimate Equation (2) separately for visitors under the Active Tab and Silent Tab.

Panel A shows that under the Active Tab, across all sales outcomes,  $\beta_1$  is positive and statistically significant (Columns (1)-(3)). The magnitudes are large: the sensitivity to high-intent visitors increases by 81% (0.013/0.016) for sales probability, 117% (0.021/0.018) for policy count, and 110% (5.194/4.740) for total premium. AI-generated demand information causes sales to become 1.8–2.2 times more sensitive to high-intent visitors. This implies that substantial information frictions exist among agents in capturing consumer purchase intent for insurance. AI information provision also significantly lowers the sensitivity of sales to low-intent visitors, suggesting that information also causes agents to reallocate time and attention across sales leads. I find no significant changes of sensitivity to the middle-intent visitor cohorts.<sup>48</sup>

<sup>47</sup>A cost of a longer window is drop of sample size. For example, when requiring a minimum of 10 weeks (*w10*) as the observational window, the total number of qualified agents is 5,369.

<sup>48</sup>*A priori*, sources of sales are not necessarily a within-agent allocation, because allocation on one consumer cannot fully determine

The coefficient in the control group ( $\beta_2$ ) demonstrates that, for example, in Column (1), when there is only human intelligence, if a visitor’s purchase intent moves from low (the omitted group) to high, the chance that she actually purchases a policy from the agent increases by 1.6%.<sup>49</sup> The positive and statistically significant coefficients in the control group ( $\beta_2$ ) also suggest that the algorithm itself is valid in predicting consumers’ actual purchase behavior.

Columns (4)-(6) in Panel A show no significant changes in sensitivity of sales to purchase intent under the Silent Tab across all outcomes. It is likely that agents almost exclusively pay attention to visitors under the Active Tab.<sup>50</sup>

Algorithmic information is displayed on the app in different formats. The intent tags in Panel A are category information. The results support Mullainathan, Schwartzstein, and Shleifer (2008)’s model of coarse thinking, where individuals “group situations into categories and apply the same model of inference to all situations within a category.” Panel B shows that agents also respond to fine-grained information – the continuous % score – which otherwise would not enter humans’ decision making and which highlights machines’ advantage. Sensitivity of sales to *Score* is significantly higher in the treatment group.<sup>51</sup> Panel C shows no significant differences in the sensitivity of sales to *Rank* between treatment and control group. Perhaps, compared to coarse category information and fine-grained information, rank of consumer purchase intent is not particularly new to the agent’s existing mental model.

**Attention Allocation.** One advantage of conducting the experiment on a mobile app is that data on agents’ app behavior allows me to uncover the *process* of their decision making. Table 4 examines how agents allocate attention to information components on the app screen using app clickstream data (within the treatment group) and Equation (3). The dependent variable is *Click* – an indicator variable that takes the value of one if the agent clicked the visitor on a day, and zero otherwise. I show that across the six columns under both Active Tab and Silent Tab, agents tend to click visitors tagged with high-intent, scored higher, and ranked on top of the screen.<sup>52</sup> The results on app clickstream are consistent with the findings on actual sales conversion (Table 3), both of which *reveal* agents’ selective attention to high-intent consumers.

**Heterogeneity by Agents’ Information Processing Capacity.** In the face of AI-generated information, a decision maker’s innate information processing capacity matters for how much she benefits from learning. In

---

that on the others. Besides, one consumer may also bring to the agent sales leads from her own social network, which is common in the insurance market. Therefore, I estimate Equation (2) without including agent fixed effects. However, Appendix Table A.2 shows that results remain similar by including agent fixed effects.

<sup>49</sup>The mean of *Any Policy* (i.e., conversion rate) in the control group is 3.6%, which means out of 100 mobile visitors who had responded to an agent’s advertisements on social media, only 3 to 4 will purchase an insurance policy eventually.

<sup>50</sup>The algorithm also displays the intent tags based on raw predicted score and the continuous % score computed as the 14-day-on-14-day change of raw predicted score, under the 14-day Active and Silent Tabs. Appendix Table A.3 reports the results using visitors under the 14-day Active and Silent Tabs and shows no significance for  $\beta_1$  except for a 10% significant increase in sensitivity for *Score* in Panel B. This suggests that agents pay less attention to 14-day Active (Silent) Tab than 7-day Active (Silent) Tab. In the main paper, I estimate Equation (2) using visitors under the 7-day Active and Silent tabs, which, unless mentioned, also applies to heterogeneity tests based on Equation (2).

<sup>51</sup>Results hold by controlling for screen *Rank* as visitors are displayed on the screen in descending order of *Score*.

<sup>52</sup>Results hold when including intent tags, *Score*, and *Rank* in one specification.

Appendix Table A.4, *College* in Panel A is an indicator variable that takes the value of one if an agent has a bachelor degree or beyond, and zero otherwise. *Experienced* in Panel B is an indicator variable that takes the value of one if an agent has above-median work experience (11 months) measured by the number of months that elapsed from the time the agent joined the agency to the time the agent entered the app page for the first time during the experiment period, and zero otherwise. *Top Performer* in Panel C is an indicator variable that takes the value of one if an agent's pre-treatment sales performance (i.e., policy count) is in the top quartile among all agents, and zero otherwise. I show that learning effects are stronger among agents who are less educated, less experienced, or underperforming (prior to the experiment). This points to the fact that agents with weak information processing capacity benefit more from improved information processing by AI.<sup>53</sup> From a Bayesian updating perspective, a decision maker's prior matters for belief update. Agents who are less educated, less experienced, or underperforming have weaker priors and are more open to new information.

**Consumer Engagement.** As an outcome of better learning about consumer purchase intent, AI information provision has the potential to facilitate more active consumer engagement and accelerate the replacement of agent-consumer relations. Appendix Table A.5 reports the information treatment effects on visitor composition, using Equation (1). New visitors visit the agent for the first time after entry (i.e., when the agent enters the app page and receives treatment). Former visitors visit the agent both before and after entry. Zombie visitors visit the agent only before entry but never again after. I find that agents in the treatment group have a higher share of new visitors and more zombie visitors. The results suggest that AI-generated demand estimates help agents form more new relational capital, re-energize vintage capital, and discard used capital.

### 7.3 Sales Composition

Table 5 examines the information treatment effects on sales composition, using Equation (1). Panel A splits the performance variables by policy term length. Short-term policy has a term length of one year or less.<sup>54</sup> Long-term policy has a term length over one year. I show that information provision improves agents' sales productivity for both short-term and long-term policies. Panel B shows that agents in the treatment group sell a 13% higher share of long-term policy (0.017/0.128) and premiums (0.026/0.199), suggesting that the algorithm is more helpful for insurance that is harder to sell. Visitors need to think more before insuring a longer term, which might result in a more complex pattern of advertising responses, thus requiring higher information processing capacity of agents. This provides evidence that the algorithm aggregates valuable consumer demand information worth learning by agents.

Appendix Table A.6 examines the information treatment effects on agents' sales performance by relationship strength, using Equation (1). Performance variables are split by whether the policyholder is a new or a former client. New vs. former is defined relative to when the agent enters the app page and receives

---

<sup>53</sup>This echoes chess grand master Garry Kasparov's view that "*Weak human + machine + better process was superior to a strong computer alone and, more remarkably, superior to a strong human + machine + inferior process.*" See this [Harvard Business Review article](#).

<sup>54</sup>Travel insurance, for example, could have a term length of less than one year.

treatment. I show that performance gains, in terms of policy count, mainly come from more sales to new clients, suggesting that the algorithm is more helpful for weak advisory relations where information asymmetry is higher.

#### 7.4 Agent Earnings and Incentives

Agent earnings come directly from the commission income tied to premiums. I examine how AI-generated demand estimates affect agent earnings and incentives. Table 6 Panel A shows that information provision improves agents' total commission income in the treatment group by 16% (91.326/579.292) and average commission income per policy by 22% (68.714/ 306.263), relative to that of the control group.

To study how agents' financial incentives might respond to accessing AI-based information on consumer demand, I then investigate the level and the structure of commission rates among all policies sold. Panel B shows that there are no significant changes in the average level of commission rates (Columns (1) and (2)). However, products sold in the treatment group have a more disperse commission rate (commission/premium) (Column (3)) and more unique commission rates (Column (4)). This points to the possibility that agents' financial incentives may respond to the information provision of consumer demand. Appendix Table A.7 tests whether information provision changes the sensitivity of agent incentives to visitors' predicted purchase intent using Equation (2). The dependent variable is an indicator variable that takes the value of one if the average commission rate of all policies sold to the visitor is above the sample median (20%) across all agent-visitors, and zero otherwise. I find that a high-intent visitor is more likely to be the policyholder of a high-commission rate product in the treatment group, as indicated by the coefficients on  $Treat \times High-intent$  in Column (1) and  $Treat \times Score$  in Column (2). The magnitude is large: the chance of selling a high-commission rate product to a high-intent visitor increases by 183% (0.011/0.006). Information provision causes agents' financial incentives to become 2.83 times more sensitive to high-intent visitors. The results are consistent with information-driven consumer discrimination.

Appendix Figure A.2 presents histograms of customers' mobile visiting patterns around policy sales for the treatment and the control group by high vs. low policy commission rate. The sample is restricted to 4,547 agents who sold at least one new policy during the experiment period, including 2,246 agents in the treatment group and 2,301 agents in the control group. Visiting dates are restricted to dates prior to the end of the experiment (November 16, 2021). The unit of observation is an agent-policy-visiting date. Day 0 is the date a policy is sold. Visiting distributions in the 100 days before and 100 days after the policy date are presented. The *y-axis* is the fraction for each bin (the sum of bar heights equals one). Bin width is 5 days. For example, Panel A shows that for the treatment (control) group, 10% (9.5%) of mobile visits occurred between day -5 and day -1 prior to selling a low-commission rate policy. A clear difference is in Panel B, which shows that consumers in the treatment group are more active in visiting agents' posts around the sale of a high-commission rate policy (day -25 to day 15) than those in the control group, suggesting that accessing AI-based demand predictions facilitates agents' consumer engagement when the monetary return from selling a policy is higher.



## 7.5 Implications for Consumers

The above findings suggest that learning about a consumer's purchase intent allows agents to profit from online advertising, presumably by reallocating sources of earnings across the consumer base. On InsurTech platforms and in digital marketing on a social network platform, the cost of advertising is low for this to happen. But are consumers exploited by more informed agents? Directly examining whether information treatment eventually improves selection of a more suitable product for consumers would require a benchmark specifying the optimal insurance choices for a given consumer (Gurun, Matvos, and Seru, 2016). I provide suggestive evidence that accessing AI-generated demand information may not realize agents' self-interests at the cost of sacrificing consumers' interests. First, Appendix Table A.8 shows that agents in the treatment group offer a more diverse set of products, suggesting expanded search efforts for more suitable products. Second, using policy cancellation as a measure of sales quality and product fit, Appendix Table A.9 finds no evidence that the information of predicted purchase intent distorts agent incentives by improving sales and personal earnings at the cost of offering poor-fitting products.

## 7.6 Information Treatment Effects on Selection

Beyond sales, insurance agents also play a critical role in collecting risk information for the insurer at the underwriting stage. In this section, I examine whether and how AI-assisted agent decisions and choices affect selection, thus shedding light on whether and how agents' incentives for risk assessment will change and reconciling those contrasting predictions proposed in Section 1. I employ empirical models of testing adverse selection in insurance markets to reveal whether AI-generated demand information induces information gain or information loss around consumer risk.

Table 7 examines selection in different agent-visitor cohorts using the risk-coverage correlation model at the policy level (Equation (4) and the model detailed in Section 6). Risk is measured as the logarithm of claimed amount.<sup>55</sup> A positive and significant correlation between risk (ex-post realization of loss) and insurance coverage is the necessary condition of adverse selection, suggesting that high risks buy more insurance. To examine the role of intermediation frictions in selection in the context of this paper, I further include insurer fixed effects and agent fixed effects one at a time in Equation (4) in Panel B and Panel C. Across the three models, I show that adverse selection exists among high- and middle-intent consumers for agents in the treatment group (Columns (2)-(4)), but not among those of the low-intent group nor in the control group (Columns (5) and (1)).<sup>56</sup> If the insurance amount per person increases by 1%, the claimed amount increases by 0.15%-0.5%. The increased  $R^2$ , especially when adding agent fixed effects, highlights the critical role of individual agents in explaining claim risk.

Table 8 examines the information treatment effects on selection using the risk-coverage correlation model fully interacted with *Treat* and predicted intent tags. The results show that AI-generated demand information

---

<sup>55</sup>Appendix Table A.10 shows robustness by measuring risk with a claim dummy.

<sup>56</sup>I also estimate Equation (4) separately for policies sold among agent-visitors in the control-high-intent group, control-middle-intent group, and control-low-intent group. None of these subgroups shows evidence for adverse selection.

leads to an increased correlation between risk and insurance coverage. The coefficients on *Log Insurance Amount*  $\times$  *Middle-intent*  $\times$  *Treat* and *Log Insurance Amount*  $\times$  *High-intent*  $\times$  *Treat* suggest that adverse selection in the treatment-middle-intent group and treatment-high-intent group is significantly more severe than it is in the respective control groups.

## 7.7 Mechanisms Underlying Adverse Selection

Increased adverse selection in the treatment group suggests that AI-generated demand information might discourage agents' own information acquisition, resulting in a crowding out of consumer risk information. I investigate a number of non-mutually exclusive explanations for why a crowding out would occur.

**Rational Inattention.** Under the theory of rational inattention, human decision makers selectively pay more attention to more important things (Sims, 2003; Bartoš et al., 2016; Maćkowiak, Matějka, and Wiederholt, 2023). They acquire information to maximize utility net of information costs and adjust attention allocation in response to changes in incentives (Dean and Neligh, 2023). Insurance agents spend costly effort acquiring customers as well as collecting consumer risk information for initial underwriting. AI demand estimations could affect agents' attention allocation between sales and risk assessment. When better informed about consumer demand for insurance, they exert more effort on sales and less effort on collecting risk information.

Results in Table 7 and Table 8 partially support this mechanism, in that adverse selection increased among high- and middle-intent consumers for agents in the treatment group, but not among those of the low-intent group. AI demand estimations not only reallocate agents' attention across their consumer base in terms of realizing sales (Section 7.2), but also reallocate agents' effort spent on sales and risk assessment.

To further support this mechanism, in Table 9 Panel A, I conduct a heterogeneity test by examining how the treatment effects on adverse selection vary with the sales "profitability" associated with a mobile visitor. *Former Client* is an indicator variable that takes the value of one if a visitor has purchased insurance from an agent before the day the agent entered the app page and saw the visitor, and zero otherwise. A former client is highly attractive to the agent due to lower cost of conversion and higher chance of second sales, especially when AI predicts her to have higher intent. The positive and significant coefficient on *Log Insurance Amount*  $\times$  *Treat*  $\times$  *Middle-intent*  $\times$  *Former Client* suggests that adverse selection worsened among former clients, compared to new clients, even when a former client is predicted to have a medium level of purchase intent by AI.<sup>57</sup>

**Saliency-driven Inattention.** Contrary to models of rational inattention, saliency-driven selective inattention suggests that it is not the information content but rather the saliency of the information that alters decision makers' attention allocation.<sup>58</sup> When AI provides demand estimates and makes sales-related signals more salient, insurance agents might simply follow the app direction and spend more (less) time on sales-related (risk assessment-related) activities.

---

<sup>57</sup>In the interest of space, only interaction terms with key-prediction coefficients are presented.

<sup>58</sup>See, for example, Bordalo, Gennaioli, and Shleifer (2013), Frydman and Wang (2020), and Medina (2021), among others.

In Table 9 Panel B, I conduct a heterogeneity test by examining how the treatment effects on adverse selection vary with the salience of the AI-generated demand predictions. *Has Recent Visiting Records* is an indicator variable that takes the value of one if a visitor has visiting records in the past 7 days, and zero otherwise. According to recency effects, more recent information and experiences have stronger influences on decision makers' working memory, probability judgment, and subsequent choices (Camerer and Loewenstein, 2011; Fredrickson and Kahneman, 1993; Schreiber and Kahneman, 2000). However, the algorithm in this paper uses a visitor's visiting records over a way longer window as the key inputs to predict a visitor's purchase intent, making human-algorithm (mis)alignment likely to occur. I posit that the scope of learning is very limited under high human-algorithm alignment when a visitor has recent visiting records and is predicted as a high-intent visitor by the algorithm. Under such a scenario, the AI prediction is less likely to reveal new information while salience is more likely to be at play. The insignificant coefficients on  $\text{Log Insurance Amount} \times \text{Treat} \times \text{Middle-intent} \times \text{Has Recent Visiting Records}$  and  $\text{Log Insurance Amount} \times \text{Treat} \times \text{High-intent} \times \text{Has Recent Visiting Records}$  do not support salience-driven inattention as the mechanism. Hence, salience alone does not explain a crowding out of agents' attention on risk assessment.

**Weak Incentives for Collecting Risk Information.** Compared to strong and direct incentives in boosting sales, agents' incentives for providing high-quality risk information to insurers might be relatively weak and indirect. However, agents do face the costs of providing low-quality risk information. For example, they could lose their relationships with insurers if they bring in too many high risks. The agency, insurers, and regulators closely monitor the quality of the underwritten pool and carry out several enforcement mechanisms. For example, the collaborating agency has created a blacklist in which agents are banned from selling insurance on the platform permanently or for a certain period of time. The sources of this list include the agency itself, insurers, and regulators. In the analysis sample, 6% of agents have ever been blacklisted.

One important reason that incentives for collecting high-quality risk information might be weak is that monetary or non-monetary punishment will materialize in the long term. Agents may fail to foresee the future costs in a multi-period setting or realize this is a repetitive game, thus crowding out their efforts spent on collecting risk information at present.

In Table 9 Panel C, I conduct a heterogeneity test by examining how the treatment effects on adverse selection vary with agents' incentives to maintain their relationships with insurers. *High Insurer Concentration (#)* is an indicator variable that takes the value of one if an agent has sold insurance from a below-sample-median number of insurers (median = 7), up until the day before the agent entered the app page and saw the visitor, and zero otherwise. I argue that if agents sell insurance products only from a small set of insurers, they have stronger incentives to provide high-quality risk information to their suppliers.<sup>59</sup> The negative and significant coefficients on  $\text{Log Insurance Amount} \times \text{Treat} \times \text{Middle-intent} \times \text{High Insurer Concentration (\#)}$  and  $\text{Log Insurance Amount} \times \text{Treat} \times \text{High-intent} \times \text{High Insurer Concentration (\#)}$  suggest that the crowding out effect is weaker among agents with stronger incentives to conduct better risk assessment so as to maintain

<sup>59</sup>A similar setting can be found in the corporate finance literature studying how customer concentration raises firms' risk exposure along the supply chain. See, for example, Campello and Gao (2017), among others.

relationships with insurers.

**AI's Substitution of Risk Information Acquisition.** Another possibility is that AI-generated demand information is substitutive to risk information collected by agents, thus reducing their own laborious information acquisition from other sources. This speaks to how AI demand estimates are able to reveal consumer risk information, if agents in practice are aware of the “correlation” between consumer demand for insurance and consumer risk type, and attempt to draw such an inference from AI demand estimates.

In Table 10, I construct a measure of the extent to which AI demand predictions mask consumer heterogeneity (e.g., risk and risk preference in [Arrow \(1972\)](#) and [Cohen and Einav \(2007\)](#)).<sup>60</sup> The idea is that some advertisements are more likely to attract high-risk consumers to purchase insurance – for example, *7 Ways Staying Up Late Could Harm Your Health* and *5 Signs of Cardiovascular Disease* – while some advertisements are more likely to raise a consumer's risk awareness, such as *Insurance 101*. When the unique number of advertisements that a visitor has visited is higher, it is harder for an agent to infer consumer risk profile deterministically. AI-generated demand information will be less substitutive to agents' own information acquisition around consumer risk, leading to less crowding out. In Panel A, *# Ads* is the unique number of advertisements that a visitor has visited over the past 180 days. In Panel B, *High # Ads* is an indicator variable that takes the value of one if the unique number of advertisements that a visitor has visited over the past 180 days is in the top quartile ( $> 6$ ) among all visitors, and zero otherwise. The negative and significant coefficients on  $\text{Log Insurance Amount} \times \text{Treat} \times \# \text{ Ads}$  and  $\text{Log Insurance Amount} \times \text{Treat} \times \text{High} \# \text{ Ads}$  suggest that the crowding out effect on human-collected information is weaker when AI demand predictions contain less clear information about consumer risk profile, supporting the information substitution mechanism.<sup>61</sup>

**Does AI Pick Lemons?** In the models of [Akerlof \(1970\)](#) and [Rothschild and Stiglitz \(1976\)](#), consumers' willingness to pay for insurance increases in their (privately known) risk type or expected costs. It is likely that the high-intent consumers predicted by AI are also high risks, if risk is the main driver of consumers' interests in responding to insurance-related advertisements. As shown in Section 7.2, with AI demand predictions, agents reallocate their efforts toward converting high-intent consumers. As a result, AI may draw agents' attention to the “lemon” market.

However, in Appendix Table A.11, I show that AI demand estimates (tag/score/rank) do not predict ex-post claim outcomes. Many factors other than consumers' own riskiness can trigger their interests in responding to insurance-related advertisements: Health problems of family members, friends, or neighbors; news media; disasters (e.g., COVID-19); risk perception/awareness; better knowledge of insurance; or some irrational responses.

---

<sup>60</sup>This is similar to [Einav et al. \(2016\)](#) in which risk scores are designed to predict an individual's expected costs in a *statistical* sense but do not concern the underlying heterogeneity about consumer health and endogenous spending response to insurance in Medicare Part D, as would be done in *economic* models.

<sup>61</sup>This measure captures the degree of *whether* AI-generated demand forecasts can help agents infer consumer risk profile or not, instead of the *level* of consumer risk. Moreover, high-intent consumers predicted by AI are not necessarily high risks (as detailed later in Appendix Table A.11). Therefore, in Table 10, I did not further interact *Treat* with intent tags.

**Moral Hazard of Consumers.** Ex-post moral hazard of consumers will strengthen the positive correlation between coverage and ex-post loss in Equation (4). In the setting of this paper, this might be less of a concern as the AI treatment is at the agent level and consumers are not aware of the on-going experiment. However, moral hazard of consumers might be *affected* by agents if AI demand predictions lead agents to spend less time doing ex-post risk management (e.g., healthcare management or prevention efforts). This conjecture is difficult to test directly and beyond the scope of this paper, but a promising area for future work, namely, examining the role of intermediary agents in reducing ex-post realization of loss. Except for moral hazard, a positive coverage-claim correlation might also be consistent with a higher risk of claiming. Better disentangling these potential channels from adverse selection will be an important part of my future work.

Taken together, the evidence suggests that a combination of mechanisms including rational inattention, weak incentives for high-quality risk assessment, and AI's substitution of information acquisition, contribute to a crowding out of risk information and increased adverse selection.

**Attention Allocation Between Consumer Demand and Risk: Evidence from App Usage.** Taking advantage of the app clickstream data to reveal agents' attention allocation between consumer demand and risk, I provide further evidence that supports a crowding out of risk information.

Table 11 examines the information treatment effects on agents' attention allocation to consumer risk information using data on app behaviors and applying Equation (1). The sample is restricted to 4,687 agents who sold at least one policy during the experiment period, including 2,321 agents in the treatment group and 2,366 agents in the control group. In Panel A, the dependent variable is the ratio of the number of times an agent viewed the health declaration pages over the total number of times an agent viewed the health declaration pages and the visitor management pages. The results show that agents in the treatment group allocate less time and attention to viewing the health declaration pages, which contain the insurer's/product's underwriting policy and the risk exposure information requested by the insurer.

Panel B examines the information treatment effects on changing the sensitivity of agents' sales to risk assessment activities on the app.<sup>62</sup> I interact *Treat* in Equation (1) with *# View Product Detail Page*, which is the number of times an agent viewed the product detail pages. Dependent variables are the number of times an agent viewed the health declaration pages (Column (1)); the average length of the input characters when the agent searched diseases, which measures the coarseness of risk information collected (Column (2)); and the number of times an agent viewed the underwriting result pages (Column (3)).<sup>63</sup> Across all three outcome variables, the negative and significant coefficients on *# View Product Detail Page*  $\times$  *Treat* show that AI demand estimates significantly lowered the sensitivity of agents' sales to risk assessment activities. Specifically, information provision lowered the sensitivity of viewing the product detail pages to viewing the health declaration pages by 64% (0.030/0.047); to the average length of the input characters when the agent searched diseases by 100% (0.001/0.001); and to viewing the underwriting result pages by 67% (0.008/0.012).

---

<sup>62</sup>This test alleviates the concern that results in Table 11 Panel A are mechanically driven by the fact that the experiment was implemented in the visitor management center of the app.

<sup>63</sup>This underwriting result is generated by the app, instead of by the insurer, to give agents a preview of potential underwriting decisions.

At the extensive margin, AI-generated demand information crowds out the amount of risk information collected by humans. At the intensive margin, AI-generated demand information might also crowd out true risk information conditional on information acquisition, if agents are subject to moral hazard and choose to hide some private and non-verifiable information regarding the expected claim cost of consumers. As these app behaviors such as viewing and disease search will not be captured by the insurers, the data used here reveals how much effort agents spent on the *process* of collecting risk information, and is not contaminated by agents' intentional avoidance of using those app functions over concerns about leaking important risk information to the insurers and potential business rejections. Therefore, the results in Table 11 also tend to support a crowding out of attention to consumer risk information at the extensive margin.

## 7.8 Product Selection as Indirect Pricing

Pricing or incentive-compatible contracts are among the many instruments used by insurance companies to mitigate the negative impacts of adverse selection (Veiga and Weyl, 2016). Although agents do not have direct pricing power, they can match riskier customers to more expensive products via product selection. Their choice set is even broader in the setting of an independent insurance agency where product menus are from multiple insurers.

Appendix Table A.12 examines the treatment effects on agents' (indirect) pricing behaviors by testing whether AI demand predictions change the risk-price correlation at the policy level. Risk is measured as the logarithm of claimed amount. Price is measured as the logarithm of premium. The insignificant coefficients on  $\text{Log Claim Amount} \times \text{Treat}$  suggest that there are no differences of risk-price correlation between agents in the treatment and control groups. Therefore, treated agents bring in riskier customers without using pricing or product selection to achieve stronger incentive compatibility (separating equilibrium).<sup>64</sup>

## 8 Conclusions

In this paper, I study how AI affects humans' attention allocation and information production in a multitasking environment by examining the impacts of AI on human intermediation of insurance contracts. I analyze a large-scale randomized field experiment conducted by a top insurance agency in China. In the experiment, the firm provided treated agents with an AI-based prediction of a consumer's demand for insurance, based on how the consumer had responded to advertising content on the largest Chinese social network platform. I show that AI demand predictions based on big data may facilitate *cherry-picking* for agents but fail to achieve *lemon-dropping* for insurers. Regarding *cherry-picking*, AI demand predictions shift agents' attention to high-intent consumers, improving agents' sales productivity by 14%. High-commission rate products are more likely to be sold to high-intent consumers, suggesting information-driven consumer discrimination. Regarding *lemon-dropping*, AI-based demand information reduces agents' own information acquisition about consumer risk and increases adverse selection, consistent with attention models and a crowding out of risk information.

---

<sup>64</sup>Results are robust by using a claim dummy to measure risk.

I highlight an important unintended consequence of AI on human-intermediated markets, namely that AI-generated information might crowd out human-collected information. Humans capitalize on AI-generated data. When AI interacts with the incentive scheme, improved information processing may deter information acquisition and exacerbate agency conflicts in a multitasking environment. In the insurance market or selection markets more generally, this side effect is reflected in information loss about consumer expected cost, which increases adverse selection and impairs market efficiency. The results may generalize to other markets with a strong selection feature where consumer demand and risk are tightly linked. One leading example will be the credit market.

The setting of this paper is in China, where independent insurance agencies set up on technology platforms are gaining momentum, compared to affiliated agents in insurance companies or independent agents in traditional brick-and-mortar intermediaries. Implications from the findings in this paper are not specific to China or to markets in developing economies with similar conditions, however. First, the application of AI studied in this paper – predicting consumer purchase intent – is common and applicable to many markets. AI helps delineate the market, but humans eventually construct the market quality. Second, in more mature insurance markets, frictions such as selective attention and conflict of interests would still exist among insurance agents. How individual decision makers in intermediaries respond to AI is the key to realizing gains from AI. In human-intermediated markets, this calls for research on how AI could be better designed and how incentive systems should cooperate on AI-assisted human decisions.

Although the large-scale field setting offers several advantages, it is important to note that my findings are based on data from one company and one AI treatment. My findings are most likely to generalize to settings where humans make effort choices in a multitasking environment and face a trade-off between the quantity and quality dimensions of their jobs. The quality dimension often requires humans to collect or to screen based on soft information. Compensation schemes may be inefficient as people are compensated too much for volume and too little for risk exposure. Risk is inherently harder to measure and longer-term than volume. This is certainly not unique to insurance sales. Other examples include mortgage originators, traders, or even CEOs. Returns to AI hinge on incentive scheme, job design, and human attention constraints. The findings about the tension between AI information processing and human information acquisition might also occur in those other settings. Moreover, as an experimental study, I showed cross-sectional/local variations to help evaluate externality. The conditions I examine for the (in)effectiveness of AI demand predictions are not exhaustive; future researchers could explore additional factors. I also supplied contextual and institutional information to assist with potential adoption in new settings.

This paper carries important implications for big data and AI in selection markets. First, in most cases in selection markets, consumer behavioral data is by nature two-dimensional as it contains information about both consumer demand and expected cost. Uni-dimensional data processing facilitates demand predictions but masks consumer heterogeneity in the drivers of demand (e.g., risk and risk preference), which are inherently hard to separate from behavioral data. Algorithms by design achieve the statistical purpose but need to take more insights from economic models. Second, data technologies are often segmented, which might be rooted in

the organizational structure (within a firm) or industry evaluation system (across firms). For example, marketing and risk assessment divisions design their own locally-powerful AI separately. When such AI interacts with the endogenous attention allocation of human users, demand predictions aimed at business expansion may have side effects on risk (cost reduction). The supply of AI requires more coordination in the production process.

Many jobs in insurance and finance resolve information asymmetries during the sales/transaction process. Intermediary agents are in a multitasking environment. Completely separating the sales and risk assessment tasks into two jobs (pure sales agent and pure underwriter) will be difficult. Risk assessment is not always a routine task and often involves soft information collection where the information collector and the decision maker (information interpreter) are often the same person (Liberti and Petersen, 2019). For example, insurance sales agents do the first part of underwriting in the field. More promising resolutions might be to design better performance measures and incentive contracts (e.g., deferral or clawback by incorporating ex-post realization of loss or profitability (Hoffmann, Inderst, and Opp, 2021, 2022), or to employ stronger monitoring and enforcement mechanisms to alter skewed attention allocation, maintain incentives for risk assessment, and reduce potential information loss.

AI technologies are evolving fast. A comparison of this paper with Brynjolfsson, Li, and Raymond (2023) on the role of generative AI highlights the diverse ways AI technologies can shape the workplace, affecting job tasks, information processing, and the overall quality of labor output. First, regarding the type of AI technology, this paper focuses on predictive AI, while theirs focuses on generative AI. Both AI aim to assist workers by distilling best practices from training big data related to customer-agent interactions. It is natural and meaningful to ask how a better or more “powerful” AI would affect workers in both settings, and indeed in other papers studying a particular AI technology. My findings are most likely to provide answers to settings where humans multitask under imperfect incentive scheme and contractual design. Second, regarding the nature of job tasks, the subjects of my paper are professional sales agents in the financial services industry, whose primary job is customer acquisition, involving collecting risk information for underwriting approval; while theirs are customer support agents in a software company, whose main job is to answer technical support questions from small business owners, i.e., problem-solving for existing customers. Third, regarding the nature of AI-generated information and the information environment, the purpose of AI is to provide best practices for searching and acquiring customers in my paper, and for servicing customers in their paper, respectively. The setting in my paper involves a hybrid mode of insurance sales – online interactions on social media and offline customer-agent interactions. AI codifies and trains the big data generated from online interactions, predicts potential customers’ purchase intent, and helps insurance agents identify sales leads. This hybrid mode provides an ideal laboratory to study how AI-generated information (i.e., demand information) affects humans’ collection of other types of information (i.e., risk information, or soft information/tacit knowledge in general). The customer support setting in their paper mainly involves online conversations and call environments, where generative AI is trained on large amounts of customer-agent conversations to identify conversational patterns that best predict call resolution and handle time. Fourth, this paper explores the challenges of multitasking in the context of AI implementation. For insurance sales agents, the presence of AI raises questions about attention allocation



and agency issues, as they navigate the dual roles of acquiring new customers and assessing their risks, which stems from the nature of selection markets. Meanwhile, customer support agents may also need to achieve a balance between resolving customer issues effectively and maintaining high levels of service quality, as gauged by the customer satisfaction metric. The trade-off between quantity and quality of workers' output broadly exists in many product and services markets. Fifth, methodologically, the paper employs a randomized field experiment to assess the short-term impacts of AI on sales agents' performance. This approach differs from the staggered introduction of the AI tool in the generative AI study, which examines the long-term effects on customer support agents' performance.

How to better integrate AI with human decision making will be an important research topic that generates more follow-up work from different fields. What dimensions of risk information are missed when incentives are distorted, for example? What if AI predicts consumer risk using the same data? How will insurance agents respond to such AI-generated risk information? What nudges could behavioral scientists implement when AI becomes more "complete" but provides information that is not aligned with human incentives? How should AI information be designed and released? What are the welfare effects of AI-generated demand predictions? Although this paper refrains from making any welfare claims with the current data, assessing the welfare effects of AI-assisted intermediary agent decisions is an important area of future work. More theories, experiments, and empirical work are needed to address these open questions.

## References

- Acemoglu, Daron, Gary Anderson, David Beede, Catherine Buffington, Eric Childress, Emin Dinlersoz, Lucia Foster, Nathan Goldschlag, John Haltiwanger, Zachary Kroff, Pascual Restrepo, and Nikolas Zolas. 2022. “Automation and the Workforce: A Firm-Level View from the 2019 Annual Business Survey.” Working Paper.
- Adams, William, Liran Einav, and Jonathan Levin. 2009. “Liquidity Constraints and Imperfect Information in Subprime Lending.” *American Economic Review*, 99 (1):49–84.
- Agarwal, Sumit, Swee H. Ang, Yongheng Deng, and Yonglin Wang. 2021. “Mortgage Brokers and the Effectiveness of Regulatory Oversight.” *Management Science*, 67 (8):5278–5300.
- Agarwal, Sumit and Itzhak Ben-David. 2018. “Loan Prospecting and the Loss of Soft Information.” *Journal of Financial Economics*, 129 (3):608–628.
- Agarwal, Sumit, Souphala Chomsisengphet, and Chunlin Liu. 2010. “The Importance of Adverse Selection in the Credit Card Market: Evidence from Randomized Trials of Credit Card Solicitations.” *Journal of Money, Credit and Banking*, 42 (4):743–754.
- Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2019. “Prediction, Judgment, and Complexity: A Theory of Decision-Making and Artificial Intelligence.” In *The Economics of Artificial Intelligence: An Agenda*, edited by Ajay Agrawal, Joshua Gans, and Avi Goldfarb. Chicago, IL: University of Chicago Press, 89–114.
- Aizawa, Naoki and You Suk Kim. 2018. “Advertising and Risk Selection in Health Insurance Markets.” *American Economic Review*, 108 (3):828–867.
- Akerlof, George A. 1970. “The Market for “Lemons”: Quality Uncertainty and the Market Mechanism.” *Quarterly Journal of Economics*, 84 (3):488–500.
- Allen, Jason, Robert Clark, Jean-François Houde, Shaoteng Li, and Anna V. Trubnikova. 2023. “The Role of Intermediaries in Selection Markets: Evidence from Mortgage Lending.” Working Paper.
- Anagol, Santosh, Shawn Cole, and Shayak Sarkar. 2017. “Understanding the Advice of Commissions-Motivated Agents: Evidence from the Indian Life Insurance Market.” *Review of Economics and Statistics*, 99 (1):1–15.
- Arrow, Kenneth J. 1972. “Essays in the Theory of Risk-Bearing.” *Journal of Political Economy*, 80 (6):1328–1329.
- Autor, David H., Frank Levy, and Richard J. Murnane. 2003. “The Skill Content of Recent Technological Change: An Empirical Exploration.” *Quarterly Journal of Economics*, 118 (4):1279–1333.
- Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson. 2024. “Artificial Intelligence, Firm Growth, and Product Innovation.” *Journal of Financial Economics*, 151:103745.
- Barbu, Alexandru. 2023. “Ex-Post Loss Sharing in Consumer Financial Markets.” Working Paper.
- Bartoš, Vojtěch, Michal Bauer, Julie Chytilová, and Filip Matějka. 2016. “Attention Discrimination: Theory and Field Experiments with Monitoring Information Acquisition.” *American Economic Review*, 106 (6):1437–1475.
- Barwick, Panle J. and Parag A. Pathak. 2015. “The Costs of Free Entry: An Empirical Study of Real Estate Agents in Greater Boston.” *RAND Journal of Economics*, 46 (1):103–145.
- Barwick, Panle J., Parag A. Pathak, and Maisy Wong. 2017. “Conflicts of Interest and Steering in Residential Brokerage.” *American Economic Journal: Applied Economics*, 9 (3):191–222.
- Bauer, Daniel, James Tyler Leverty, Joan Schmit, and Justin Sydnor. 2021. “Symposium on Insure-tech, Digitalization, and Big-data Techniques in Risk Management and Insurance.” *Journal of Risk and Insurance*, 88 (3):525–528.
- Bauhoff, Sebastian. 2012. “Do Health Plans Risk-Select? An Audit Study on Germany’s Social Health Insurance.” *Journal of Public Economics*, 96 (9-10):750–759.
- Berg, Tobias, Valentin Burg, Ana Gombović, and Manju Puri. 2020. “On the Rise of FinTechs: Credit Scoring Using Digital Footprints.” *Review of Financial Studies*, 33 (7):2845–2897.
- Berg, Tobias, Andreas Fuster, and Manju Puri. 2022. “Fintech Lending.” *Annual Review of Financial*

- Economics, 14 (1):187–207.
- Berg, Tobias, Manju Puri, and Jörg Rocholl. 2020. “Loan Officer Incentives, Internal Rating Models, and Default Rates.” Review of Finance, 24 (3):529–578.
- Bergstresser, Daniel, John M. R. Chalmers, and Peter Tufano. 2009. “Assessing the Costs and Benefits of Brokers in the Mutual Fund Industry.” Review of Financial Studies, 22 (10):4129–4156.
- Bertrand, Marianne, Dean Karlan, Sendhil Mullainathan, Eldar Shafir, and Jonathan Zinman. 2010. “What’s Advertising Content Worth? Evidence from a Consumer Credit Marketing Field Experiment.” Quarterly Journal of Economics, 125 (1):263–306.
- Bonelli, Maxime. 2023. “Data-driven Investors.” Working Paper.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2013. “Salience and Consumer Choice.” Journal of Political Economy, 121 (5):803–843.
- Brown, Jason, Mark Duggan, Ilyana Kuziemko, and William Woolston. 2014. “How Does Risk Selection Respond to Risk Adjustment? New Evidence from the Medicare Advantage Program.” American Economic Review, 104 (10):3335–3364.
- Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond. 2023. “Generative AI at Work.” Working Paper.
- Buchak, Greg, Gregor Matvos, Tomasz Piskorski, and Amit Seru. 2022. “Why is Intermediating Houses So Difficult? Evidence from iBuyers.” Working Paper.
- Bundorf, M. Kate, Maria Polyakova, and Ming Tai-Seale. 2024. “How Do Consumers Interact with Digital Expert Advice? Experimental Evidence from Health Insurance.” Management Science. Forthcoming.
- Burbidge, John B., Lonnie Magee, and A. Leslie Robb. 1988. “Alternative Transformations to Handle Extreme Values of the Dependent Variable.” Journal of the American Statistical Association, 83 (401):123–127.
- Calvino, Flavio and Luca Fontanelli. 2023. “A Portrait of AI Adopters Across Countries: Firm Characteristics, Assets’ Complementarities and Productivity.” Working Paper.
- Camerer, Colin F. and George Loewenstein. 2011. “Behavioral Economics: Past, Present, Future.” In Advances in Behavioral Economics, edited by Colin F. Camerer, George Loewenstein, and Matthew Rabin. Princeton, NJ: Princeton University Press, 3–51.
- Campello, Murillo and Janet Gao. 2017. “Customer Concentration and Loan Contract Terms.” Journal of Financial Economics, 123 (1):108–136.
- Cao, Sean S., Wei Jiang, Junbo L. Wang, and Baozhong Yang. 2023. “From Man vs. Machine to Man + Machine: The Art and AI of Stock Analyses.” Journal of Financial Economics. Forthcoming.
- Caplin, Andrew, David J. Deming, Søren Leth-Petersen, and Ben Weidmann. 2023. “Allocative Skill.” Working Paper.
- Carey, Colleen. 2017. “Technological Change and Risk Adjustment: Benefit Design Incentives in Medicare Part D.” American Economic Journal: Economic Policy, 9 (1):38–73.
- Carrell, Scott E. and James E. West. 2010. “Does Professor Quality Matter? Evidence from Random Assignment of Students to Professors.” Journal of Political Economy, 118 (3):409–432.
- Carson, James M., Randy E. Dumm, and Robert E. Hoyt. 2007. “Incentive Compensation and the Use of Contingent Commissions: The Case of Smaller Distribution Channel Members.” Journal of Insurance Regulation, 25 (4):53–67.
- Cebul, Randall D., James B. Rebitzer, Lowell J. Taylor, and Mark E. Votruba. 2011. “Unhealthy Insurance Markets: Search Frictions and the Cost and Quality of Health Insurance.” American Economic Review, 101 (5):1842–1871.
- Chak, Ida, Karen Croxson, Francesco D’Acunto, Jonathan Reuter, and Jonathan Shaw. 2022. “Improving Household Debt Management with Robo-advice.” Working Paper.
- Cheng, Jiang, Elyas Elyasiani, and Tzu-Ting Lin. 2010. “Market Reaction to Regulatory Action in the Insurance Industry: The Case of Contingent Commission.” Journal of Risk and Insurance, 77 (2):347–368.
- Chiappori, Pierre-Andre and Bernard Salanie. 2000. “Testing for Asymmetric Information in Insurance Markets.” Journal of Political Economy, 108 (1):56–78.
- Christoffersen, Susan E. K., Richard Evans, and David K. Musto. 2013. “What Do Consumers’ Fund Flows

- Maximize? Evidence from Their Brokers' Incentives." Journal of Finance, 68 (1):201–235.
- Cohen, Alma and Liran Einav. 2007. "Estimating Risk Preferences from Deductible Choice." American Economic Review, 97 (3):745–788.
- Cohen, Alma and Peter Siegelman. 2010. "Testing for Adverse Selection in Insurance Markets." Journal of Risk and Insurance, 77 (1):39–84.
- Cole, Shawn, Martin Kanz, and Leora Klapper. 2015. "Incentivizing Calculated Risk-Taking: Evidence from an Experiment with Commercial Bank Loan Officers." Journal of Finance, 70 (2):537–575.
- Coleman, Braiden, Kenneth Merkley, and Joseph Pacelli. 2022. "Human Versus Machine: A Comparison of Robo-Analyst and Traditional Research Analyst Investment Recommendations." The Accounting Review, 97 (5):221–244.
- Cooper, Robert W. 2007. "Spitzer's Allegations of the Anticompetitive Effects of Contingent Commissions: A Shot Truly Heard Around the World." Journal of Insurance Regulation, 26 (1):83–113.
- Cummins, J. D., Neil Doherty, Gerald Ray, and Terri Vaughan. 2006. "The Insurance Brokerage Industry Post-October 2004." Risk Management and Insurance Review, 9 (2):89–108.
- Cummins, J. David and Neil A. Doherty. 2006. "The Economics of Insurance Intermediaries." Journal of Risk and Insurance, 73 (3):359–396.
- Cupach, William R. and James M. Carson. 2002. "The Influence of Compensation on Product Recommendations Made by Insurance Agents." Journal of Business Ethics, 40 (2):167–176.
- Cutler, David M. and Richard Zeckhauser. 2004. "Extending the Theory to Meet the Practice of Insurance." Brookings-Wharton Papers on Financial Services, 2004 (1):1–53.
- D'Acunto, Francesco, Pulak Ghosh, Rajiv Jain, and Alberto G. Rossi. 2023. "How Costly are Cultural Biases? Evidence from FinTech." Working Paper.
- D'Acunto, Francesco, Nagpurnanand Prabhala, and Alberto G. Rossi. 2019. "The Promises and Pitfalls of Robo-Advising." Review of Financial Studies, 32 (5):1983–2020.
- D'Acunto, Francesco and Alberto G. Rossi. 2023a. "IT Meets Finance: Financial Decision Making in the Digital Era." In Handbook of Financial Decision Making, edited by Gilles Hilary and David McLean. Cheltenham, United Kingdom: Edward Elgar Publishing Ltd., 336–354.
- . 2023b. "New Frontiers of Robo-Advising: Consumption, Saving, Debt Management, and Taxes." In Machine Learning and Data Sciences for Financial Markets: A Guide to Contemporary Practices, edited by Agostino Capponi and Charles-Albert Lehalle. Cambridge, United Kingdom: Cambridge University Press, 9–32.
- D'Acunto, Francesco, Alberto G. Rossi, and Michael Weber. 2023. "Crowdsourcing Peer Information to Change Spending Behavior." Journal of Financial Economics. Forthcoming.
- de Mel, Suresh, Craig McIntosh, Ketki Sheth, and Christopher Woodruff. 2022. "Can Mobile-Linked Bank Accounts Bolster Savings? Evidence from a Randomized Controlled Trial in Sri Lanka." Review of Economics and Statistics, 104 (2):306–320.
- Dean, Mark and Nathaniel Neligh. 2023. "Experimental Tests of Rational Inattention." Journal of Political Economy, 131 (12):3415–3461.
- Decarolis, Francesco and Andrea Guglielmo. 2017. "Insurers' Response to Selection Risk: Evidence from Medicare Enrollment Reforms." Journal of Health Economics, 56:383–396.
- DeFusco, Anthony A., Huan Tang, and Constantine Yannelis. 2022. "Measuring the Welfare Cost of Asymmetric Information in Consumer Credit Markets." Journal of Financial Economics, 146 (3):821–840.
- Deming, David J. 2021. "The Growing Importance of Decision-Making on the Job." Working Paper.
- Dionne, Georges and Casey Rothschild. 2014. "Economic Effects of Risk Classification Bans." Geneva Risk and Insurance Review, 39 (2):184–221.
- Dizon-Ross, Rebecca. 2019. "Parents' Beliefs about Their Children's Academic Ability: Implications for Educational Investments." American Economic Review, 109 (8):2728–2765.
- Dobbie, Will, Andres Liberman, Daniel Paravisini, and Vikram Pathania. 2021. "Measuring Bias in Consumer Lending." Review of Economic Studies, 88 (6):2799–2832.

- Dugast, Jérôme and Thierry Foucault. 2018. “Data Abundance and Asset Price Informativeness.” Journal of Financial Economics, 130 (2):367–391.
- . 2022. “Equilibrium Data Mining and Data Abundance.” Working Paper.
- Dumm, Randy E. and Robert E. Hoyt. 2003. “Insurance Distribution Channels: Markets in Transition.” Journal of Insurance Regulation, 22 (1):27–47.
- Egan, Mark. 2019. “Brokers versus Retail Investors: Conflicting Interests and Dominated Products.” Journal of Finance, 74 (3):1217–1260.
- Egan, Mark, Shan Ge, and Johnny Tang. 2022. “Conflicting Interests and the Effect of Fiduciary Duty: Evidence from Variable Annuities.” Review of Financial Studies, 35 (12):5334–5386.
- Egan, Mark, Gregor Matvos, and Amit Seru. 2019. “The Market for Financial Adviser Misconduct.” Journal of Political Economy, 127 (1):233–295.
- Einav, Liran and Amy Finkelstein. 2011. “Selection in Insurance Markets: Theory and Empirics in Pictures.” Journal of Economic Perspectives, 25 (1):115–138.
- Einav, Liran, Amy Finkelstein, Raymond Kluender, and Paul Schrimpf. 2016. “Beyond Statistics: The Economic Content of Risk Scores.” American Economic Journal: Applied Economics, 8 (2):195–224.
- Einav, Liran, Amy Finkelstein, and Jonathan Levin. 2010. “Beyond Testing: Empirical Models of Insurance Markets.” Annual Review of Economics, 2 (1):311–336.
- Einav, Liran, Amy Finkelstein, and Neale Mahoney. 2021. “The IO of Selection Markets.” In Handbook of Industrial Organization, vol. 5, edited by Kate Ho, Ali Hortaçsu, and Alessandro Lizzeri. Amsterdam: Elsevier B.V., 389–426.
- Einav, Liran, Mark Jenkins, and Jonathan Levin. 2012. “Contract Pricing in Consumer Credit Markets.” Econometrica, 80 (4):1387–1432.
- Eling, Martin, Ruo Jia, and Yi Yao. 2017. “Between-Group Adverse Selection: Evidence from Group Critical Illness Insurance.” Journal of Risk and Insurance, 84 (2):771–809.
- Eling, Martin, Davide Nuessle, and Julian Staubli. 2022. “The Impact of Artificial Intelligence Along the Insurance Value Chain and on the Insurability of Risks.” Geneva Papers on Risk and Insurance-Issues and Practice, 47 (2):205–241.
- Fang, Hanming, Michael P. Keane, and Dan Silverman. 2008. “Sources of Advantageous Selection: Evidence from the Medigap Insurance Market.” Journal of Political Economy, 116 (2):303–350.
- Farboodi, Maryam and Laura Veldkamp. 2020. “Long-Run Growth of Financial Data Technology.” American Economic Review, 110 (8):2485–2523.
- Finkelstein, Amy and Kathleen McGarry. 2006. “Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market.” American Economic Review, 96 (4):938–958.
- Foerster, Stephen, Juhani T. Linnainmaa, Brian T. Melzer, and Alessandro Previtero. 2017. “Retail Financial Advice: Does One Size Fit All?” Journal of Finance, 72 (4):1441–1482.
- Fredrickson, Barbara L. and Daniel Kahneman. 1993. “Duration Neglect in Retrospective Evaluations of Affective Episodes.” Journal of Personality and Social Psychology, 65 (1):45–55.
- Frydman, Cary and Baolian Wang. 2020. “The Impact of Salience on Investor Behavior: Evidence from a Natural Experiment.” Journal of Finance, 75 (1):229–276.
- Gabaix, Xavier. 2014. “A Sparsity-Based Model of Bounded Rationality.” Quarterly Journal of Economics, 129 (4):1661–1710.
- Gargano, Antonio, Marco Giacoletti, and Elvis Jarnecic. 2023. “Local Experiences, Search, and Spillovers in the Housing Market.” Journal of Finance, 78 (2):1015–1053.
- Gargano, Antonio and Alberto G. Rossi. 2018. “Does It Pay to Pay Attention?” Review of Financial Studies, 31 (12):4595–4649.
- . 2022. “Goal Setting and Saving in the FinTech Era.” Journal of Finance. Forthcoming.
- Geruso, Michael and Timothy Layton. 2020. “Upcoding: Evidence from Medicare on Squishy Risk Adjustment.” Journal of Political Economy, 128 (3):984–1026.
- Ghosh, Chinmoy and James I. Hilliard. 2012. “The Value of Contingent Commissions in the Property-Casualty

- Insurance Industry: Evidence from Stock Market Returns.” Journal of Risk and Insurance, 79 (1):165–192.
- Gilbukh, Sonia and Paul S. Goldsmith-Pinkham. 2023. “Heterogeneous Real Estate Agents and the Housing Cycle.” Working Paper.
- Goldstein, Itay and Liyan Yang. 2017. “Information Disclosure in Financial Markets.” Annual Review of Financial Economics, 9 (1):101–125.
- Gravelle, Hugh, Matt Sutton, and Ada Ma. 2010. “Doctor Behaviour under a Pay for Performance Contract: Treating, Cheating and Case Finding.” The Economic Journal, 120 (542):F129–F156.
- Gruber, Jonathan, Benjamin R. Handel, Samuel H. Kina, and Jonathan T. Kolstad. 2021. “Managing Intelligence: Skilled Experts and Decision Support in Markets for Complex Products.” Working Paper.
- Grullon, Gustavo, George Kanatas, and James P. Weston. 2004. “Advertising, Breadth of Ownership, and Liquidity.” Review of Financial Studies, 17 (2):439–461.
- Gupta, Arpit and Christopher Hansman. 2022. “Selection, Leverage, and Default in the Mortgage Market.” Review of Financial Studies, 35 (2):720–770.
- Gurun, Umit G. and Alexander W. Butler. 2012. “Don’t Believe the Hype: Local Media Slant, Local Advertising, and Firm Value.” Journal of Finance, 67 (2):561–598.
- Gurun, Umit G., Gregor Matvos, and Amit Seru. 2016. “Advertising Expensive Mortgages.” Journal of Finance, 71 (5):2371–2416.
- Hastings, Justine, Ali Hortaçsu, and Chad Syverson. 2017. “Sales Force and Competition in Financial Product Markets: The Case of Mexico’s Social Security Privatization.” Econometrica, 85 (6):1723–1761.
- He, Zhiguo, Sheila Jiang, Douglas Xu, and Xiao Yin. 2023. “Investing in Lending Technology: IT Spending in Banking.” Working Paper.
- Healy, Paul M. 1985. “The Effect of Bonus Schemes on Accounting Decisions.” Journal of Accounting and Economics, 7 (1):85–107.
- Heider, Florian and Roman Inderst. 2012. “Loan Prospecting.” Review of Financial Studies, 25 (8):2381–2415.
- Hendel, Igal, Aviv Nevo, and Francois Ortalo-Magne. 2009. “The Relative Performance of Real Estate Marketing Platforms: MLS versus FSBOMadison.com.” American Economic Review, 99 (5):1878–1898.
- Hertzberg, Andrew, Jose M. Liberti, and Daniel Paravisini. 2010. “Information and Incentives Inside the Firm: Evidence from Loan Officer Rotation.” Journal of Finance, 65 (3):795–828.
- Hilliard, James I., Laureen Regan, and Sharon Tennyson. 2013. “Insurance Distribution.” In Handbook of Insurance, edited by Georges Dionne. New York, NY: Springer, 689–727.
- Hoffman, Mitchell, Lisa B. Kahn, and Danielle Li. 2018. “Discretion in Hiring.” Quarterly Journal of Economics, 133 (2):765–800.
- Hoffmann, Florian, Roman Inderst, and Marcus Opp. 2021. “Only Time Will Tell: A Theory of Deferred Compensation.” Review of Economic Studies, 88 (3):1253–1278.
- . 2022. “The Economics of Deferral and Clawback Requirements.” Journal of Finance, 77 (4):2423–2470.
- Holmstrom, Bengt and Paul Milgrom. 1991. “Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design.” Journal of Law, Economics, and Organization, 7:24–52.
- Honka, Elisabeth, Ali Hortaçsu, and Maria A. Vitorino. 2017. “Advertising, Consumer Awareness, and Choice: Evidence from the U.S. Banking Industry.” RAND Journal of Economics, 48 (3):611–646.
- Hoyt, Robert E., James M. Carson, and Randy E. Dumm. 2005. “An Examination of the Role of Insurance Producers and Compensation in the Insurance Industry.” Independent Insurance Agents and Brokers Association. Alexandria, VA.
- Hu, Zhongchen. 2022. “Social Interactions and Households’ Flood Insurance Decisions.” Journal of Financial Economics, 144 (2):414–432.
- Jansen, Mark, Hieu Quang Nguyen, and Amin Shams. 2023. “Rise of the Machines: The Impact of Automated Underwriting.” Management Science. Forthcoming.
- Kacperczyk, Marcin, Stijn Van Nieuwerburgh, and Laura Veldkamp. 2014. “Time-Varying Fund Manager Skill.” Journal of Finance, 69 (4):1455–1484.


- . 2016. “A Rational Theory of Mutual Funds’ Attention Allocation.” *Econometrica*, 84 (2):571–626.
- Karaca-Mandic, Pinar, Roger Feldman, and Peter Graven. 2018. “The Role of Agents and Brokers in the Market for Health Insurance.” *Journal of Risk and Insurance*, 85 (1):7–34.
- Karlan, Dean and Jonathan Zinman. 2009. “Observing Unobservables: Identifying Information Asymmetries with a Consumer Credit Field Experiment.” *Econometrica*, 77 (6):1993–2008.
- Keys, Benjamin J., Tanmoy Mukherjee, Amit Seru, and Vikrant Vig. 2010. “Did Securitization Lead to Lax Screening? Evidence from Subprime Loans.” *Quarterly Journal of Economics*, 125 (1):307–362.
- Keys, Benjamin J., Amit Seru, and Vikrant Vig. 2012. “Lender Screening and the Role of Securitization: Evidence from Prime and Subprime Mortgage Markets.” *Review of Financial Studies*, 25 (7):2071–2108.
- Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. 2018. “Human Decisions and Machine Predictions.” *Quarterly Journal of Economics*, 133 (1):237–293.
- Koijen, Ralph S. J., Stijn Van Nieuwerburgh, and Motohiro Yogo. 2016. “Health and Mortality Delta: Assessing the Welfare Cost of Household Insurance Choice.” *Journal of Finance*, 71 (2):957–1010.
- Koijen, Ralph S. J. and Motohiro Yogo. 2022. “New Perspectives on Insurance.” *Review of Financial Studies*, 35 (12):5275–5286.
- Lee, Sung Kwan. 2020. “Fintech Nudges: Overspending Messages and Personal Finance Management.” Working Paper.
- Levitt, Steven D. and Chad Syverson. 2008. “Market Distortions When Agents Are Better Informed: The Value of Information in Real Estate Transactions.” *Review of Economics and Statistics*, 90 (4):599–611.
- Liberti, José M. and Mitchell A. Petersen. 2019. “Information: Hard and Soft.” *Review of Corporate Finance Studies*, 8 (1):1–41.
- Liebenberg, Andre P., James M. Carson, and Randy E. Dumm. 2012. “A Dynamic Analysis of the Demand for Life Insurance.” *Journal of Risk and Insurance*, 79 (3):619–644.
- Linnainmaa, Juhani T., Brian T. Melzer, and Alessandro Previtero. 2021. “The Misguided Beliefs of Financial Advisors.” *Journal of Finance*, 76 (2):587–621.
- Liu, Lei, Guangli Lu, and Wei Xiong. 2024. “The Big Tech Lending Model.” Working Paper.
- Lou, Dong. 2014. “Attracting Investor Attention through Advertising.” *Review of Financial Studies*, 27 (6):1797–1829.
- Ludwig, Jens and Sendhil Mullainathan. 2021. “Fragile Algorithms and Fallible Decision-Makers: Lessons from the Justice System.” *Journal of Economic Perspectives*, 35 (4):71–96.
- Lyonnet, Victor and Léa H. Stern. 2022. “Venture Capital (Mis)Allocation in the Age of AI.” Working Paper.
- Ma, Yu-Luen, Nat Pope, and Xiaoying Xie. 2014. “Contingent Commissions, Insurance Intermediaries, and Insurer Performance.” *Risk Management and Insurance Review*, 17 (1):61–81.
- Maćkowiak, Bartosz, Filip Matějka, and Mirko Wiederholt. 2023. “Rational Inattention: A Review.” *Journal of Economic Literature*, 61 (1):226–273.
- Medina, Paolina C. 2021. “Side Effects of Nudging: Evidence from a Randomized Intervention in the Credit Card Market.” *Review of Financial Studies*, 34 (5):2580–2607.
- Mehrotra, Ateev, Sonya Grier, and R. A. Dudley. 2006. “The Relationship Between Health Plan Advertising and Market Incentives: Evidence of Risk-Selective Behavior.” *Health Affairs*, 25 (3):759–765.
- Misra, Sanjog. 2019. “Selling and Sales Management.” In *Handbook of the Economics of Marketing*, vol. 1, edited by Jean-Pierre H. Dube and Peter E. Rossi. Amsterdam: Elsevier B.V., 441–496.
- Mullainathan, Sendhil, Antoinette Schoar, and Markus Noeth. 2012. “The Market for Financial Advice: An Audit Study.” Working Paper.
- Mullainathan, Sendhil, Joshua Schwartzstein, and Andrei Shleifer. 2008. “Coarse Thinking and Persuasion.” *Quarterly Journal of Economics*, 123 (2):577–619.
- Neuman, Patricia, Ed Maibach, Katharine Dusenbury, Michelle Kitchman, and Pam Zupp. 1998. “Marketing HMOs to Medicare Beneficiaries.” *Health Affairs*, 17 (4):132–139.
- Newhouse, Joseph P., Mary Price, J. M. McWilliams, John Hsu, and Thomas G. McGuire. 2015. “How Much Favorable Selection Is Left in Medicare Advantage?” *American Journal of Health Economics*, 1 (1):1–26.

- Outreville, J. François. 2013. “The Relationship Between Insurance and Economic Development: 85 Empirical Papers for a Review of the Literature.” Risk Management and Insurance Review, 16 (1):71–122.
- Oyer, Paul. 1998. “Fiscal Year Ends and Nonlinear Incentive Contracts: The Effect on Business Seasonality.” Quarterly Journal of Economics, 113 (1):149–185.
- Qian, Jun (QJ), Philip E. Strahan, and Zhishu Yang. 2015. “The Impact of Incentives and Communication Costs on Information Production and Use: Evidence from Bank Lending.” Journal of Finance, 70 (4):1457–1493.
- Regan, Laureen and Sharon Tennyson. 1996. “Agent Discretion and the Choice of Insurance Marketing System.” Journal of Law and Economics, 39 (2):637–666.
- Rejda, George E. and Michael J. McNamara. 2014. “Insurance Company Operations.” In Principles of Risk Management and Insurance 12th Edition, edited by George E. Rejda and Michael J. McNamara. New Jersey, NJ: Pearson, 106.
- Reuter, Jonathan and Eric Zitzewitz. 2006. “Do Ads Influence Editors? Advertising and Bias in the Financial Media.” Quarterly Journal of Economics, 121 (1):197–227.
- Richter, Andreas, Jörg Schiller, and Harris Schlesinger. 2014. “Behavioral Insurance: Theory and Experiments.” Journal of Risk and Uncertainty, 48 (2):85–96.
- Robles-Garcia, Claudia. 2022. “Competition and Incentives in Mortgage Markets: The Role of Brokers.” Working Paper.
- Rossi, Alberto G. and Stephen Utkus. 2021. “Who Benefits from Robo-advising? Evidence from Machine Learning.” Working Paper.
- Rothschild, Michael and Joseph Stiglitz. 1976. “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information.” Quarterly Journal of Economics, 90 (4):629–649.
- Schiller, Jörg. 2008. “Matching and Risk Classification in Insurance Markets with Intermediation.” Working Paper.
- Schreiber, Charles A. and Daniel Kahneman. 2000. “Determinants of the Remembered Utility of Aversive Sounds.” Journal of Experimental Psychology: General, 129 (1):27–42.
- Shapiro, Bradley T. 2020. “Advertising in Health Insurance Markets.” Marketing Science, 39 (3):587–611.
- Sims, Christopher A. 2003. “Implications of Rational Inattention.” Journal of Monetary Economics, 50 (3):665–690.
- Spatt, Chester S. 2020. “Conflicts of Interest in Asset Management and Advising.” Annual Review of Financial Economics, 12:217–235.
- Stigler, George J. 1961. “The Economics of Information.” Journal of Political Economy, 69 (3):213–225.
- Stroebel, Johannes. 2016. “Asymmetric Information about Collateral Values.” Journal of Finance, 71 (3):1071–1112.
- Sydnor, Justin. 2010. “(Over)insuring Modest Risks.” American Economic Journal: Applied Economics, 2 (4):177–199.
- Tong, Joy Tianjiao. 2022. “Advising Sophisticated Customers: Evidence from Health Insurance Brokers.” Working Paper.
- Veiga, André and E. Glen Weyl. 2016. “Product Design in Selection Markets.” Quarterly Journal of Economics, 131 (2):1007–1056.
- Zolas, Nikolas, Zachary Kroff, Erik Brynjolfsson, Kristina McElheran, David Beede, Catherine Buffington, Nathan Goldschlag, Lucia Foster, and Emin Dinlersoz. 2020. “Advanced Technologies Adoption and Use by U.S. Firms: Evidence from the Annual Business Survey.” Working Paper.

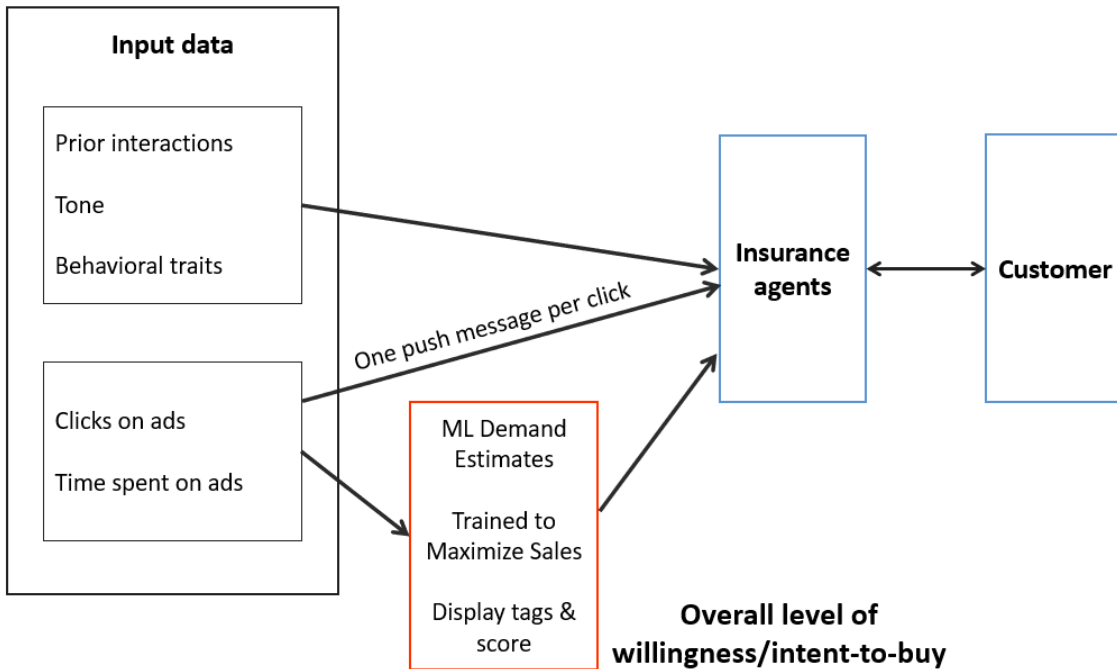


Figure 1: Agents' Inputs for Decision Making: Control vs. Treatment Group

Panel A: Control Group




Panel B: Treatment Group



**Note:** This figure presents agents' inputs for decision making in the control group (Panel A) and the treatment group (Panel B). In the experiment, the firm provided treated agents with an AI-based prediction of a consumer's purchase intent for insurance, based on how the consumer had responded to advertising content on the largest Chinese social network platform, WeChat. Agents in both the treatment and the control group had access to information about consumer response to advertising – the same information that AI is using to generate the demand estimates. Whenever a WeChat user clicks an advertisement, the agent will receive a push notification on WeChat. The agent sees which advertisement and how much time the visitor has spent on reading it. AI-generated data is only available to agents in the treatment group and displayed on their app screen. In addition, agents' inputs for decision making include personal and interactive elements (e.g., prior interactions, tone, and consumer behavioral traits) that do not enter the algorithm.

Figure 2: **Organizational Structure of the Cooperating Insurance Agency**



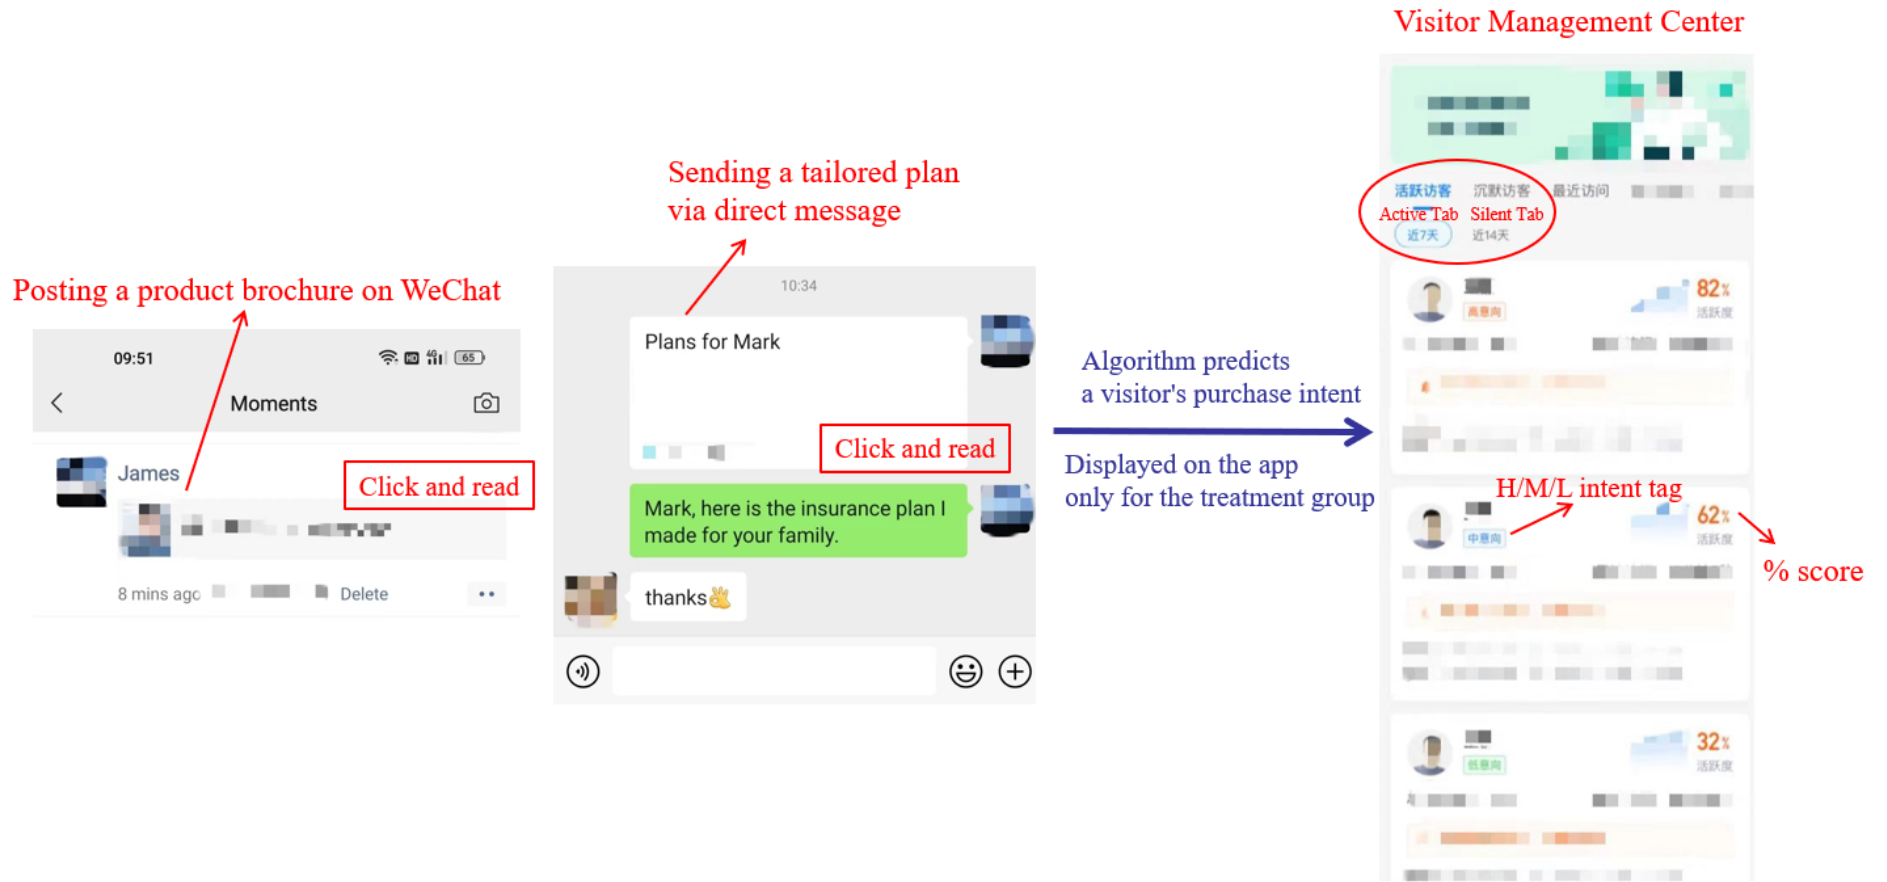
**Note:** This figure shows the organizational structure of the cooperating insurance agency.

Figure 3: **Sales of Insurance: A Multitasking Environment**

Sales	Risk Assessment
<b>information of demand – AI</b>	information of risk
quantity	quality
hard information	soft information
incentivized/rewarded	nonincentivized/nonrewarded
observable to the principal	observable to the principal
–	costs of poor risk assessment

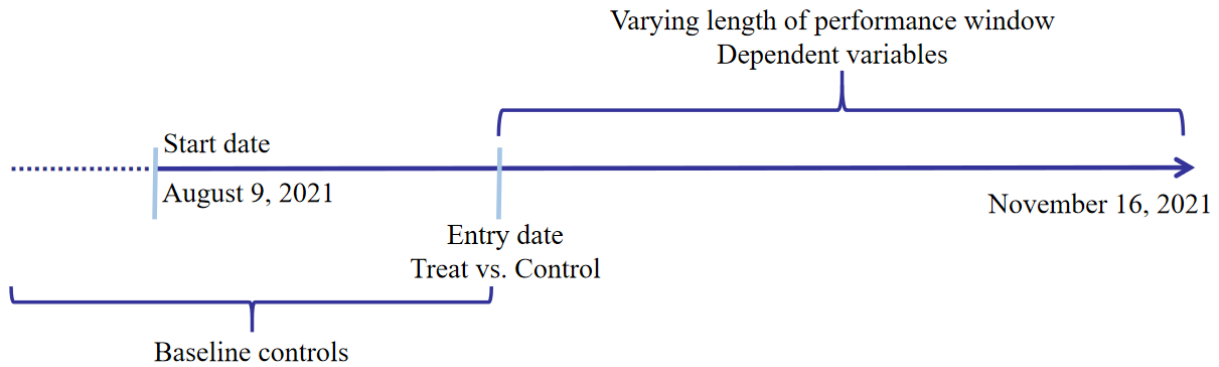
**Note:** This figure summarizes the key elements of insurance sales as a multitasking environment.

Figure 4: Examples of Advertising, Visiting, and App Page for Experiment



**Note:** This figure shows examples of how an agent advertises on WeChat by posting content or sending direct messages, how a WeChat user becomes a visitor by clicking and reading the shared content, how the app page is set up for the experiment, and what information components are displayed on the mobile screen. The pictures are pixelated as per the company's request.

Figure 5: **Timeline of Experiment and Variable Construction**



**Note:** This figure shows the timeline of the experiment and the timeframe of variable construction. The experiment is conducted by the agency from August 9 to November 16, 2021. Agents are randomly assigned to a treatment or a control group from the first time they enter the app page during this period. Dependent variables are measured from the day agents enter the app page until November 16 (or a day earlier if they leave the platform before November 16). I require at least four weeks (28 days) as the observational window; thus agents in the baseline sample enter the app page between August 9 and October 20, including 5,430 agents in the treatment group and 5,695 agents in the control group. Baseline control variables include total number of policies, total premium up until the day before agents entered the app page, and the number of days of the observational window.

Table 1: Balance Checks

	All (1)	Treatment (2)	Control (3)	Mean Diff. (4)	T-stat (5)
	N=11,125	N=5,430	N=5,695		
<b>Demographics &amp; Experience</b>					
Age	40.55	40.46	40.63	-0.17	-0.80
Female	0.51	0.51	0.51	-0.00	-0.30
College	0.30	0.30	0.29	0.01	1.49
Branch in First-tier City	0.43	0.43	0.42	0.01	0.68
APP Experience (months)	27.14	27.08	27.20	-0.12	-0.31
Work Experience (months)	10.63	10.57	10.68	-0.11	-0.58
<b>Pre-treatment Sales Performance &amp; Composition</b>					
Any Policy	0.53	0.54	0.53	0.01	1.04
# Policy	10.01	10.32	9.71	0.61	0.89
Total Premium	10265.28	10478.63	10061.86	416.77	0.47
Avg. Premium Per Policy	1417.86	1483.60	1354.01	129.59	0.64
# Policy - Short-term	8.49	8.72	8.27	0.45	0.71
# Policy - Long-term	1.51	1.59	1.44	0.16	1.17
Total Premium - Short-term	2697.30	2783.95	2614.69	169.26	0.93
Total Premium - Long-term	7567.98	7694.68	7447.18	247.50	0.30
Avg. Premium Per Policy - Short-term	392.74	399.31	386.34	12.97	1.34
Avg. Premium Per Policy - Long-term	6445.39	6285.50	6608.90	-323.40	-0.32
% Policy - Long-term	0.16	0.16	0.15	0.01	1.13
% Premium - Long-term	0.28	0.28	0.27	0.01	1.38
# Policy - New Clients	6.77	6.96	6.58	0.38	0.94
# Policy - Former Clients	3.24	3.35	3.13	0.23	0.58
Total Premium - New Clients	7357.16	7529.56	7192.79	336.77	0.49
Total Premium - Former Clients	2908.12	2949.07	2869.07	80.00	0.24
Avg. Premium Per Policy - New Clients	1391.55	1480.66	1305.21	175.46	0.87
Avg. Premium Per Policy - Former Clients	1655.18	1461.07	1849.42	-388.35	-1.02
% Policy - New Clients	0.85	0.85	0.85	-0.00	-0.76
% Premium - New Clients	0.84	0.84	0.85	-0.00	-0.67
<b>Commission Income &amp; Rate</b>					
Total Commission	3049.60	3104.98	2996.79	108.19	0.49
Avg. Commission Per Policy	446.29	464.73	428.38	36.35	1.07
Avg. Commission Rate	0.24	0.24	0.24	-0.00	-0.48
Med. Commission Rate	0.24	0.24	0.24	-0.00	-0.69
SD. Commission Rate	0.08	0.08	0.08	-0.00	-0.39
Unique # Commission Rate	4.57	4.66	4.49	0.18	1.26

Table 1: **Balance Checks**—*Continued*

	All (1)	Treatment (2)	Control (3)	Mean Diff. (4)	T-stat (5)
	N=11,125	N=5,430	N=5,695		
<b>Claims</b>					
# Claims	1.56	1.60	1.52	0.08	0.34
Claimed Amount	2122.00	2024.26	2218.37	-194.11	-0.41
Claim Ratio	0.08	0.08	0.08	-0.00	-0.02
<b>Policy Cancellation</b>					
# Canceled Policy	0.27	0.26	0.27	-0.01	-0.25
Cancellation Ratio	0.03	0.03	0.03	-0.00	-0.38
<b>Product Variety</b>					
Unique # Products	5.90	6.03	5.77	0.26	1.46
Product HHI	0.50	0.49	0.50	-0.01	-1.61
<b>App Usage</b>					
# View Health Declaration Page	13.54	14.15	12.92	1.23	1.06
# View Visitor Management Page	21.00	21.20	20.84	0.36	0.24
% View Health Declaration Page	0.46	0.46	0.46	-0.00	-0.11
# Avg. Length of Disease Search Input Characters	3.14	3.13	3.16	-0.03	-0.68
# View Underwriting Result Page	21.54	22.62	20.45	2.17	1.54

**Note:** This table presents means of pre-treatment characteristics for agents in the full baseline sample, treatment group, and control group, the mean difference between the two groups, and the t-stat of that difference. All variables are measured as of the day prior to the agent entering the experiment app page for the first time during the experiment period. *Branch in First-tier City* is an indicator variable taking the value of one if the agent is affiliated with a branch in Beijing, Shanghai, Guangzhou, or Shenzhen, and zero otherwise. *APP Experience* is the number of months since app registration. *Work Experience* is the number of months since the agent registered at the insurance agency. *Short-term* policy has a term length of one year or less. *Long-term* policy has a term length over one year. *% Policy - Long-term* is the number of long-term policies over the number of all policies sold. *% Premium - Long-term* is the total premium of long-term policies over the total premium of all policies sold. *Commission Rate* for each policy is the ratio of commission income over premium. *Claim Ratio* is the share of policies that ever had any claims. *Cancellation Ratio* is the share of policies that had been canceled. *Product HHI* is the Herfindahl-Hirschman Index (HHI) of the number of policies sold for each product. *% View Health Declaration Page* is the ratio of the number of times an agent viewed the health declaration pages over the total number of times an agent viewed the health declaration pages and the visitor management pages. Pre-treatment performance is based on historical sales to mobile visitors who have visiting records on agents' posted/shared content on WeChat. Premium, commission income, and claimed amount are in RMB. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.

**Table 2: Information Treatment Effects on Sales Performance**

	# Policy		Tot. Premium		Avg. Premium	
	(1)	(2)	(3)	(4)	(5)	(6)
Treat	0.326** (0.152)	0.311** (0.134)	300.068** (138.003)	275.733** (123.602)	120.411 (89.023)	121.344 (86.471)
Observations	11,125	11,125	11,125	11,125	4,687	4,687
Baseline Controls	No	Yes	No	Yes	No	Yes
R-squared	0.000	0.223	0.000	0.203	0.000	0.061
Control Mean	2.940	2.940	1965.433	1965.433	1081.500	1081.500

**Note:** This table reports the information treatment effects on agents' sales performance, using Equation (1). The experiment is conducted by the agency from August 9 to November 16, 2021. Agents are randomly assigned to a treatment or a control group from the first time they enter the app page during this period. Dependent variables are measured from the day agents enter the app page until November 16 (or a day earlier if they leave the platform before November 16). I require at least four weeks (28 days) as the observational window; thus agents in the baseline sample enter the app page between August 9 and October 20, including 5,430 agents in the treatment group and 5,695 agents in the control group. Dependent variables are winsorized at the 1st and 99th percentiles. *Avg. Premium* (average premium per policy) in Columns (5) and (6) are available only for 4,687 agents who sold at least one policy during the experiment period, including 2,321 agents in the treatment group and 2,366 agents in the control group. Columns (2), (4) and (6) include baseline control variables, including total number of policies, total premium up until the day before agents entered the app page, and the number of days of the observational window. Robust standard errors are used. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.

**Table 3: Information Treatment Effects on Sensitivity of Sales to Visitors' Predicted Purchase Intent**

Panel A: Sensitivity of Sales to Predicted Intent Tag

	Active Tab			Silent Tab		
	Any Policy (1)	# Policy (2)	Tot. Premium (3)	Any Policy (4)	# Policy (5)	Tot. Premium (6)
Treat × High-intent	0.013** (0.006)	0.021** (0.009)	5.194** (2.638)	0.001 (0.003)	0.001 (0.003)	0.349 (0.636)
Treat × Middle-intent	0.009 (0.006)	0.015* (0.009)	4.128 (2.570)	0.000 (0.002)	0.000 (0.002)	-0.046 (0.465)
High-intent	0.016*** (0.006)	0.018** (0.008)	4.740** (2.354)	0.016*** (0.002)	0.016*** (0.002)	3.216*** (0.470)
Middle-intent	0.004 (0.005)	0.003 (0.008)	-0.060 (2.292)	0.006*** (0.002)	0.006*** (0.002)	1.081*** (0.346)
Treat	-0.017*** (0.006)	-0.025*** (0.009)	-6.127** (2.548)	-0.002 (0.002)	-0.002 (0.002)	-0.466 (0.337)
Observations	79,563	79,563	79,563	136,861	136,861	136,861
R-squared	0.026	0.029	0.027	0.030	0.030	0.029
Control Mean	0.036	0.051	13.234	0.019	0.019	3.996

Panel B: Sensitivity of Sales to Predicted Score

	Active Tab			Silent Tab		
	Any Policy (1)	# Policy (2)	Tot. Premium (3)	Any Policy (4)	# Policy (5)	Tot. Premium (6)
Treat × Score	0.002** (0.001)	0.002** (0.001)	0.522* (0.279)	-0.001 (0.006)	-0.001 (0.006)	0.086 (1.244)
Score	0.000 (0.000)	0.001 (0.001)	0.011 (0.184)	0.032*** (0.005)	0.032*** (0.005)	6.207*** (0.945)
Treat	-0.009*** (0.003)	-0.012*** (0.004)	-2.641** (1.106)	-0.002 (0.001)	-0.002 (0.001)	-0.412 (0.288)
Observations	79,563	79,563	79,563	136,861	136,861	136,861
R-squared	0.024	0.028	0.026	0.029	0.029	0.028
Control Mean	0.036	0.051	13.234	0.019	0.019	3.996

Panel C: Sensitivity of Sales to Screen Rank

	Active Tab			Silent Tab		
	Any Policy (1)	# Policy (2)	Tot. Premium (3)	Any Policy (4)	# Policy (5)	Tot. Premium (6)
Treat × Rank	0.0000 (0.0000)	0.0000 (0.0001)	0.0141 (0.0186)	0.0000 (0.0000)	0.0000 (0.0000)	0.0045 (0.0055)
Rank	-0.0004*** (0.0000)	-0.0006*** (0.0000)	-0.1751*** (0.0139)	-0.0002*** (0.0000)	-0.0002*** (0.0000)	-0.0386*** (0.0038)
Treat	-0.0050** (0.0025)	-0.0067* (0.0036)	-1.2940 (1.1140)	-0.0022 (0.0014)	-0.0022 (0.0014)	-0.4866 (0.3062)
Observations	79,563	79,563	79,563	136,861	136,861	136,861
R-squared	0.029	0.032	0.031	0.030	0.030	0.029
Control Mean	0.036	0.051	13.234	0.019	0.019	3.996

**Note:** This table tests whether information provision changes the sensitivity of sales performance to visitors' predicted purchase intent by comparing the baseline sensitivity for the control group (coefficients on *High-intent* and *Middle-intent*) with the increase in sensitivity for the treatment group (coefficients on *Treat × High-intent* and *Treat × Middle-intent*). The coefficient on *Treat* represents the effect on sensitivity change for the omitted *Low-intent* group. The predictive algorithm displays two information components on the app: the high-, middle-, low-intent tags based on raw predicted score (Panel A) and a continuous % score computed as the 7-day-on-7-day change of the raw predicted score representing the trend of purchase intent (Panel B). Visitors with a positive (negative) % score are displayed under the Active (Silent) Tab. Screen rank (Panel C) is not an output of the algorithm but a feature of the screen display where visitors are displayed in descending order of the continuous % score. The unit of observation is an agent-visitor. For each visitor, dependent variables are measured from the day the agent sees the visitor for the first time (i.e., when the agent enters the app page and the visitor is displayed on the screen on that day) until November 16 (or a day earlier if the agent leaves the platform before November 16). Dependent variables (except for *Any Policy*) are winsorized at the 1st and 99th percentiles. Independent variables *High-intent*, *Middle-intent*, *Score*, and *Rank* are the visitor's baseline information seen by the agent on the first day. *Score* is displayed on the app screen as a % and divided by 100 in regressions for presenting coefficients. *Rank* is capped at 100. All models include baseline control variables at the agent-visitor level, including total number of policies, total premium up until the day before the agent entered the app page and saw the visitor, and the number of days of the observational window for each agent-visitor. Standard errors are clustered at the agent level. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.



Table 4: Attention to Information Components on Mobile Screen within Treatment Group

	Active Tab			Silent Tab		
	(1)	(2)	(3)	(4)	(5)	(6)
High-intent	0.0316*** (0.0050)			0.0009*** (0.0003)		
Middle-intent	0.0047 (0.0042)			0.0012*** (0.0003)		
Score		0.0093*** (0.0010)			0.0015** (0.0007)	
Rank			-0.0004*** (0.0001)			-0.0000*** (0.0000)
Observations	75,690	75,690	75,690	275,114	275,114	275,114
Agent FE	Yes	Yes	Yes	Yes	Yes	Yes
Day of Week FE	Yes	Yes	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.170	0.173	0.169	0.132	0.132	0.132
Outcome Mean	0.0310	0.0310	0.0310	0.0016	0.0016	0.0016

**Note:** This table tests how agents allocate attention to information components on the mobile screen using click data. The sample is restricted to agents in the treatment group. The predictive algorithm displays two information components on the app: the high-, middle-, low-intent tags based on a raw predicted score, and a continuous % score computed as the 7-day-on-7-day change of the raw predicted score representing the trend of purchase intent. Visitors with a positive (negative) % score are displayed under the Active (Silent) Tab. Screen rank is not an output of the algorithm but a feature of the screen display where visitors are displayed in descending order of the continuous % score. The unit of observation is an agent-day-visitor and includes all days an agent entered the app page during the experiment period. The dependent variable is an indicator variable that takes the value of one if the agent clicked a visitor on a day, and zero otherwise. Independent variables *High-intent*, *Middle-intent*, *Score*, and *Rank* are the visitor's information on each day. *Score* is displayed on the app screen as a % and divided by 100 in regressions for presenting coefficients. *Rank* is capped at 100. All models include baseline control variables at the agent-day-visitor level, including total number of policies and total premium up until the day before the agent entered the app page and saw the visitor. All models include agent, day of week, and month fixed effects. Standard errors are clustered at the agent level. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.

Table 5: **Information Treatment Effects on Sales Composition: By Policy Term Length**

Panel A: By Policy Term Length

	Short-term			Long-term		
	# Policy (1)	Tot. Premium (2)	Avg. Premium (3)	# Policy (4)	Tot. Premium (5)	Avg. Premium (6)
Treat	0.239* (0.131)	53.167* (29.780)	2.757 (9.625)	0.037*** (0.013)	234.485** (109.376)	-377.352 (863.809)
Observations	11,125	11,125	4,320	11,125	11,125	1,227
R-squared	0.197	0.187	0.008	0.217	0.158	0.031
Control Mean	2.702	640.219	310.297	0.209	1214.577	7836.177

Panel B: Sales Composition

	% Long-term Policy (1)	% Long-term Policy Premium (2)
	Treat	0.017** (0.008)
Observations	4,687	4,687
R-squared	0.067	0.103
Control Mean	0.128	0.199

**Note:** This table reports the information treatment effects on agents' sales composition by policy term length, using Equation (1). The sample includes 5,430 agents in the treatment group and 5,695 agents in the control group. *Short-term* policy has a term length of one year or less. *Long-term* policy has a term length over one year. Panel A estimates the baseline results separately by measuring performance using only short-term or long-term policies. *Avg. Premium* (average premium per policy) in Columns (3) and (6) are available only for agents who sold at least one short-term (4,320 agents) or long-term (1,227 agents) policy during the experiment period. Dependent variables are winsorized at the 1st and 99th percentiles. Panel B uses the share of long-term policies and total premiums as dependent variables. The sample includes 4,687 agents who sold at least one policy during the experiment period, including 2,321 agents in the treatment group and 2,366 agents in the control group. All models include baseline control variables, including total number of policies, total premium up until the day before agents entered the app page, and the number of days of the observational window. Robust standard errors are used. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.

Table 6: **Information Treatment Effects on Agent Earnings and Incentives**

Panel A: Commission Income

	Tot. Commission (1)	Avg. Commission (2)
Treat	91.326** (38.547)	68.714** (28.371)
Observations	11,125	4,687
R-squared	0.179	0.061
Control Mean	579.292	306.263

Panel B: Commission Rate

	Avg. Commission Rate (1)	Med. Commission Rate (2)	SD. Commission Rate (3)	Unique # Commission Rate (4)
Treat	-0.002 (0.002)	-0.004 (0.003)	0.003* (0.002)	0.146** (0.061)
Observations	4,687	4,687	3,827	4,687
R-squared	0.009	0.012	0.052	0.201
Control Mean	0.204	0.205	0.065	2.751

**Note:** This table reports the information treatment effects on agent earnings and incentives, using Equation (1). Panel A looks at total commission income and average commission income per policy. Panel B looks at commission rate features (commission/premium), including average commission rate, median commission rate, standard deviation of commission rates, and the unique number of commission rates across all policies. The sample includes 5,430 agents in the treatment group and 5,695 agents in the control group for *Tot. Commission*. The sample includes 4,687 agents who sold at least one policy during the experiment period, including 2,321 agents in the treatment group and 2,366 agents in the control group, for other dependent variables. *Tot. Commission* and *Avg. Commission* are winsorized at the 1st and 99th percentiles. All models include baseline control variables, including total number of policies, total premium up until the day before agents entered the app page, and the number of days of the observational window. Robust standard errors are used. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.

Table 7: Information Treatment Effects on Selection: Separate Estimations

Panel A: Without Insurer and Agent FE

	DV: Log Claim Amount				
	Control-All (1)	Treat-All (2)	Treat-High (3)	Treat-Middle (4)	Treat-Low (5)
Log Insurance Amount	-0.000 (0.047)	0.154*** (0.056)	0.215** (0.094)	0.182*** (0.064)	-0.141 (0.185)
Observations	3,186	2,832	1,065	1,546	214
Premium Rate	Yes	Yes	Yes	Yes	Yes
Age/Gender/Location	Yes	Yes	Yes	Yes	Yes
Insurance Type FE	Yes	Yes	Yes	Yes	Yes
Applicant-Insurant Relation FE	Yes	Yes	Yes	Yes	Yes
Insurer FE	No	No	No	No	No
Agent FE	No	No	No	No	No
R-squared	0.030	0.030	0.052	0.033	0.101
Outcome Mean	0.257	0.275	0.238	0.283	0.413

Panel B: Including Insurer FE

	DV: Log Claim Amount				
	Control-All (1)	Treat-All (2)	Treat-High (3)	Treat-Middle (4)	Treat-Low (5)
Log Insurance Amount	-0.023 (0.066)	0.292** (0.115)	0.351** (0.162)	0.416** (0.196)	-0.136 (0.202)
Observations	3,184	2,831	1,064	1,546	211
Premium Rate	Yes	Yes	Yes	Yes	Yes
Age/Gender/Location	Yes	Yes	Yes	Yes	Yes
Insurance Type FE	Yes	Yes	Yes	Yes	Yes
Applicant-Insurant Relation FE	Yes	Yes	Yes	Yes	Yes
Insurer FE	Yes	Yes	Yes	Yes	Yes
Agent FE	No	No	No	No	No
R-squared	0.039	0.048	0.118	0.059	0.131
Outcome Mean	0.257	0.275	0.238	0.283	0.380

Panel C: Including Insurer and Agent FE

	DV: Log Claim Amount				
	Control-All (1)	Treat-All (2)	Treat-High (3)	Treat-Middle (4)	Treat-Low (5)
Log Insurance Amount	0.002 (0.112)	0.427*** (0.160)	0.458* (0.250)	0.504* (0.269)	-0.241 (0.186)
Observations	2,969	2,625	923	1,396	185
Premium Rate	Yes	Yes	Yes	Yes	Yes
Age/Gender/Location	Yes	Yes	Yes	Yes	Yes
Insurance Type FE	Yes	Yes	Yes	Yes	Yes
Applicant-Insurant Relation FE	Yes	Yes	Yes	Yes	Yes
Insurer FE	Yes	Yes	Yes	Yes	Yes
Agent FE	Yes	Yes	Yes	Yes	Yes
R-squared	0.196	0.272	0.248	0.276	0.652
Outcome Mean	0.250	0.277	0.244	0.283	0.315

**Note:** This table examines selection in different agent-visitor cohorts using the risk-coverage correlation model. The analysis is at the policy level. A positive and significant correlation between risk and coverage is the necessary condition of adverse selection, suggesting that high risks buy more insurance. Risk is measured as the logarithm of claimed amount. Coverage is measured as the logarithm of insurance amount. Claim data is available only for certain products in student safety insurance, accident insurance, and health insurance. Insurance Type refers to those product categories. I accessed claim data on June 6, 2022. The sample includes all policies sold to visitors during the experiment period where all variables used in regressions are available. *Treat-High*, *Treat-Middle*, and *Treat-Low* are based on a visitor's baseline purchase intent classification seen by the agent on the first day. Standard errors are clustered at the agent level. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.

Table 8: **Information Treatment Effects on Selection: Interaction Approach**

DV: Log Claim Amount

	(1)	(2)	(3)
Log Insurance Amount	0.015 (0.034)	0.087 (0.065)	0.147 (0.121)
Log Insurance Amount × Middle-intent	0.049 (0.046)	0.030 (0.046)	0.016 (0.072)
Log Insurance Amount × High-intent	0.000 (0.041)	-0.024 (0.048)	-0.081 (0.071)
Log Insurance Amount × Middle-intent × Treat	0.235** (0.106)	0.229** (0.104)	0.104 (0.125)
Log Insurance Amount × High-intent × Treat	0.221** (0.106)	0.263** (0.113)	0.261** (0.126)
Log Insurance Amount × Treat	-0.184** (0.092)	-0.207** (0.096)	-0.097 (0.124)
Treat	2.534** (1.248)	2.796** (1.315)	
Middle-intent	-0.549 (0.576)	-0.324 (0.594)	-0.149 (0.927)
High-intent	0.043 (0.531)	0.325 (0.639)	1.018 (0.907)
Treat × Middle-intent	-3.145** (1.398)	-3.043** (1.394)	-1.644 (1.673)
Treat × High-intent	-3.141** (1.437)	-3.601** (1.524)	-3.512** (1.685)
Observations	6,026	6,024	5,604
Premium Rate	Yes	Yes	Yes
Age/Gender/Location	Yes	Yes	Yes
Insurance Type FE	Yes	Yes	Yes
Applicant-Insurant Relation FE	Yes	Yes	Yes
Insurer FE	No	No	Yes
Agent FE	No	Yes	Yes
R-squared	0.027	0.036	0.220
Outcome Mean	0.265	0.265	0.263

**Note:** This table examines the information treatment effects on selection using the risk-coverage correlation model fully interacted with *Treat* and predicted intent tags. The analysis is at the policy level. Risk is measured as the logarithm of claimed amount. Coverage is measured as the logarithm of insurance amount. The sample includes all policies sold to visitors during the experiment period where all variables used in regressions are available. *High-intent* and *Middle-intent* (*Low-intent* is omitted) are based on a visitor's baseline purchase intent classification seen by the agent on the first day. Standard errors are clustered at the agent level. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.

Table 9: **Information Treatment Effects on Selection: Mechanisms**

DV: Log Claim Amount

	(1)	(2)	(3)
<b>Panel A: Rational Inattention</b>			
Log Insurance Amount × Treat × Middle-intent × Former Client	0.511** (0.251)		
Log Insurance Amount × Treat × High-intent × Former Client	0.248 (0.184)		
<b>Panel B: Salience-driven Inattention</b>			
Log Insurance Amount × Treat × Middle-intent × Has Recent Visiting Records		1.452 (1.181)	
Log Insurance Amount × Treat × High-intent × Has Recent Visiting Records		1.566 (1.203)	
<b>Panel C: Weak Incentives for Collecting Risk Information</b>			
Log Insurance Amount × Treat × Middle-intent × High Insurer Concentration (#)			-0.558** (0.260)
Log Insurance Amount × Treat × High-intent × High Insurer Concentration (#)			-0.889*** (0.241)
Observations	5,604	5,604	5,604
Premium Rate	Yes	Yes	Yes
Age/Gender/Location	Yes	Yes	Yes
Insurance Type FE	Yes	Yes	Yes
Applicant-Insurant Relation FE	Yes	Yes	Yes
Insurer FE	Yes	Yes	Yes
Agent FE	Yes	Yes	Yes
R-squared	0.221	0.223	0.222
Outcome Mean	0.263	0.263	0.263

**Note:** This table examines the mechanisms underlying the information treatment effects on selection. I fully interact the risk-coverage correlation model with *Treat*, predicted intent tags, and variables of interest for heterogeneity tests. In the interest of space, only interaction terms with key-prediction coefficients are presented. The analysis is at the policy level. The sample includes all policies sold to visitors during the experiment period where all variables used in regressions are available. *High-intent* and *Middle-intent* (*Low-intent* is omitted) are based on a visitor’s baseline purchase intent classification seen by the agent on the first day. In Panel A, *Former Client* is an indicator variable that takes the value of one if a visitor has purchased insurance from an agent before the day the agent entered the app page and saw the visitor, and zero otherwise. In Panel B, *Has Recent Visiting Records* is an indicator variable that takes the value of one if a visitor has visiting records in the past 7 days, and zero otherwise. In Panel C, *High Insurer Concentration (#)* is an indicator variable that takes the value of one if an agent has sold insurance from a below-sample-median number of insurers, up until the day before the agent entered the app page and saw the visitor, and zero otherwise. Standard errors are clustered at the agent level. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.

Table 10: **Information Treatment Effects on Selection: AI's Substitution of Risk Information Acquisition**

DV: Log Claim Amount

	(1)	(2)
<b>Panel A: Continuous</b>		
Log Insurance Amount $\times$ Treat $\times$ # Ads	-0.025** (0.011)	
<b>Panel B: Dummy</b>		
Log Insurance Amount $\times$ Treat $\times$ High # Ads		-0.283* (0.161)
Observations	5,604	5,604
Premium Rate	Yes	Yes
Age/Gender/Location	Yes	Yes
Insurance Type FE	Yes	Yes
Applicant-Insurant Relation FE	Yes	Yes
Insurer FE	Yes	Yes
Agent FE	Yes	Yes
R-squared	0.218	0.220
Outcome Mean	0.263	0.263

**Note:** This table examines whether AI-generated demand information is substitutive to risk information collected by agents, thus reducing their own information acquisition. I fully interact the risk-coverage correlation model with *Treat* and a measure of the extent to which AI demand predictions mask consumer heterogeneity (e.g., risk and risk preference). When the unique number of advertisements that a visitor has visited is higher, it is harder for an agent to infer consumer risk profile deterministically. AI-generated demand information will be less substitutive to agents' own information acquisition around consumer risk, leading to less crowding out. The analysis is at the policy level. The sample includes all policies sold to visitors during the experiment period where all variables used in regressions are available. In the interest of space, only interaction terms with key-prediction coefficients are presented. In Panel A, *# Ads* is the unique number of advertisements that a visitor has visited over the past 180 days. In Panel B, *High # Ads* is an indicator variable that takes the value of one if the unique number of advertisements that a visitor has visited over the past 180 days is in the top quartile among all visitors, and zero otherwise. Standard errors are clustered at the agent level. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.

**Table 11: Information Treatment Effects on Agents' Attention Allocation to Consumer Risk Information: Evidence from App Behaviors**

Panel A: Attention to Risk Information

	% View Health Declaration Page (1)
Treat	-0.019** (0.009)
Observations	4,687
R-squared	0.030
Control Mean	0.263

Panel B: Sensitivity of Sales to Risk Assessment Activities

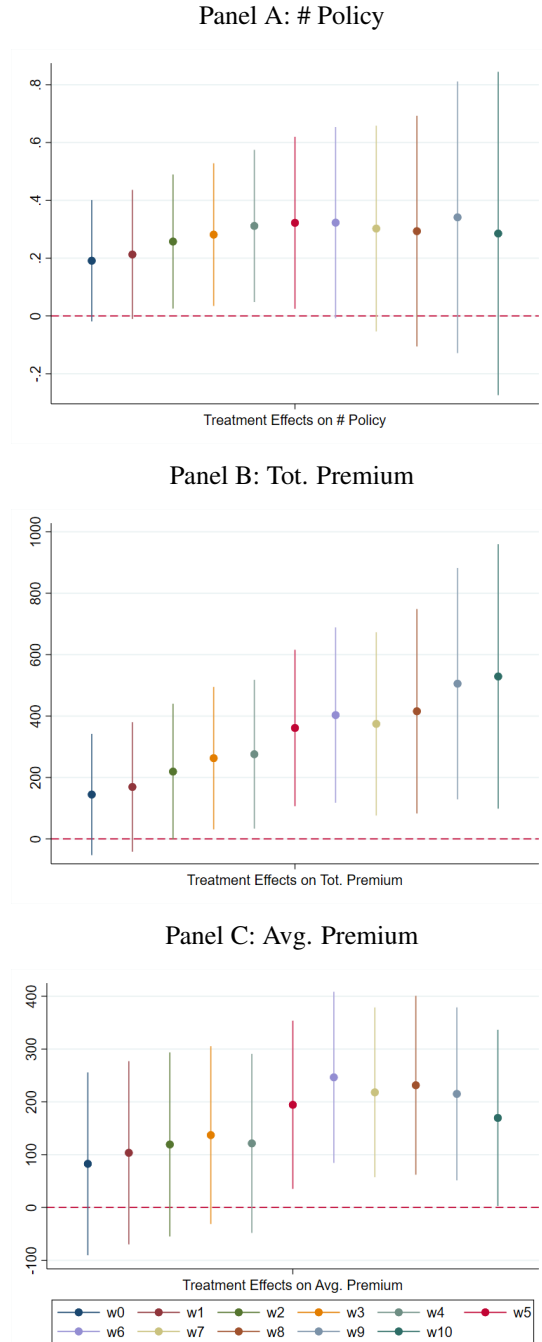
	# View Health Declaration Page (1)	Avg. Length of Disease Search Input Characters (2)	# View Underwriting Result Page (3)
# View Product Detail Page	0.047*** (0.015)	0.001*** (0.000)	0.012*** (0.002)
# View Product Detail Page × Treat	-0.030* (0.016)	-0.001*** (0.000)	-0.008*** (0.003)
Treat	7.835** (3.544)	0.277*** (0.083)	2.359*** (0.691)
Observations	4,687	4,687	4,687
R-squared	0.211	0.075	0.108
Control Mean	9.245	0.728	2.762

**Note:** This table examines the information treatment effects on agents' attention allocation to consumer risk information using data on app behaviors and applying Equation (1). The sample is restricted to 4,687 agents who sold at least one policy during the experiment period, including 2,321 agents in the treatment group and 2,366 agents in the control group. In Panel A, the dependent variable is the ratio of the number of times an agent viewed the health declaration pages over the total number of times an agent viewed the health declaration pages and the visitor management pages. Panel B examines the information treatment effects on changing the sensitivity of agents' sales to risk assessment activities on the app. I interact *Treat* in Equation (1) with *# View Product Detail Page*, which is the number of times an agent viewed the product detail pages. Dependent variables are the number of times an agent viewed the health declaration pages (Column (1)); the average length of the input characters when the agent searched diseases (Column (2)); and the number of times an agent viewed the underwriting result pages (Column (3)). All models include baseline control variables, including total number of policies, total premium up until the day before agents entered the app page, and the number of days of the observational window. Robust standard errors are used. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.



## Appendix A: Additional Figures and Tables for Online Publication

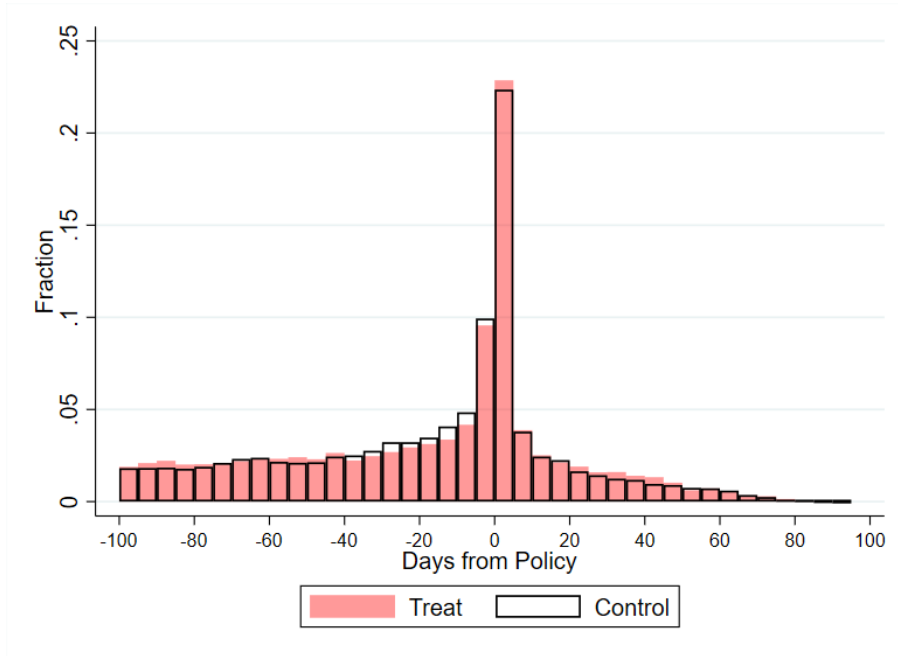
Figure A.1: Information Treatment Effects on Sales Performance: Varying Sample By Timing of Entry




**Note:** This figure reports the information treatment effects on agents' sales performance when varying the agent sample by timing of entry, using Equation (1). The experiment is conducted by the agency from August 9 to November 16, 2021. Agents are randomly assigned to a treatment or a control group from the first time they enter the app page during this period. Dependent variables are measured from the day agents enter the app page until November 16 (or a day earlier if they leave the platform before November 16). The point estimates represent coefficients of *Treat* (with their 95% confidence intervals) by requiring a different minimum number of weeks as the observational window; thus the longer the window, the earlier the agent entered the app page. For example, *w1* shows the effect when requiring at least one week for observing performance. Dependent variables are winsorized at the 1st and 99th percentiles. All models include baseline control variables, including total number of policies, total premium up until the day before agents entered the app page, and the number of days of the observational window. Robust standard errors are used.

Figure A.2: Customers' Mobile Visiting Patterns Around Policy Sales: By Policy Commission Rate

Panel A: Low Commission Rate Policy



Panel B: High Commission Rate Policy



**Note:** This figure presents histograms of customers' mobile visiting patterns around policy sales for the treatment and control groups by policy commission rate. The sample is restricted to 4,547 agents who sold at least one new policy during the experiment period, including 2,246 agents in the treatment group and 2,301 agents in the control group. Visiting dates are restricted to those prior to the end of the experiment (November 16, 2021). The unit of observation is an agent-policy-visiting date. Day 0 is the date when a policy is sold. Visiting distributions in 100 days before and 100 days after the policy dates are presented. The *y-axis* is the fraction for each bin (the sum of bar heights equals one). Bin width is 5 days. For example, Panel A shows that for the treatment (control) group, 10% (9.5%) of mobile visits occurred between day -5 and day -1 prior to selling a low-commission rate policy. Low (Panel A) and high (Panel B) commission rates are based on the median cut-off among all new policies.

Table A.1: **Information Treatment Effects on Sales Performance: Outcomes in Logs and IHS**

Panel A: Outcome in Logs

	Log # Policy		Log Tot. Premium		Log Avg. Premium	
	(1)	(2)	(3)	(4)	(5)	(6)
Treat	0.040** (0.019)	0.040** (0.017)	0.141** (0.070)	0.144** (0.065)	0.074* (0.041)	0.073* (0.039)
Observations	11,125	11,125	11,125	11,125	4,687	4,687
Baseline Controls	No	Yes	No	Yes	No	Yes
R-squared	0.000	0.217	0.000	0.152	0.001	0.085
Control Mean	0.675	0.675	2.946	2.946	5.810	5.810

Panel B: Outcome in IHS Transformation

	IHS # Policy		IHS Tot. Premium		IHS Avg. Premium	
	(1)	(2)	(3)	(4)	(5)	(6)
Treat	0.049** (0.024)	0.049** (0.021)	0.150* (0.077)	0.152** (0.071)	0.073* (0.041)	0.072* (0.040)
Observations	11,125	11,125	11,125	11,125	4,687	4,687
Baseline Controls	No	Yes	No	Yes	No	Yes
R-squared	0.000	0.209	0.000	0.148	0.001	0.084
Control Mean	0.849	0.849	3.233	3.233	6.497	6.497

**Note:** This table reports the information treatment effects on agents' sales performance, using Equation (1). Panel A uses performance outcomes in logs. Panel B uses an inverse hyperbolic sine (IHS) transformation to account for the skewed non-normal distribution of sales. The experiment is conducted by the agency from August 9 to November 16, 2021. Agents are randomly assigned to a treatment or a control group from the first time they enter the app page during this period. Dependent variables are measured from the day agents enter the app page until November 16 (or a day earlier if they leave the platform before November 16). I require at least four weeks (28 days) as the observational window; thus agents in the baseline sample enter the app page between August 9 and October 20, including 5,430 agents in the treatment group and 5,695 agents in the control group. *Avg. Premium* (average premium per policy) in Columns (5) and (6) are available only for 4,687 agents who sold at least one policy during the experiment period, including 2,321 agents in the treatment group and 2,366 agents in the control group. Columns (2), (4) and (6) include baseline control variables, including total number of policies, total premium up until the day before agents entered the app page, and the number of days of the observational window. Robust standard errors are used. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.

**Table A.2: Information Treatment Effects on Sensitivity of Sales to Visitors' Predicted Purchase Intent: Including Agent Fixed Effects**

Panel A: Sensitivity of Sales to Predicted Intent Tag

	Active Tab			Silent Tab		
	Any Policy (1)	# Policy (2)	Tot. Premium (3)	Any Policy (4)	# Policy (5)	Tot. Premium (6)
Treat × High-intent	0.015* (0.009)	0.028** (0.013)	7.171* (3.880)	0.000 (0.003)	0.000 (0.003)	0.430 (0.745)
Treat × Middle-intent	0.010 (0.008)	0.020 (0.012)	5.027 (3.680)	-0.001 (0.003)	-0.001 (0.003)	-0.094 (0.618)
High-intent	-0.008 (0.008)	-0.019 (0.012)	-5.377 (3.419)	0.016*** (0.002)	0.016*** (0.002)	3.088*** (0.533)
Middle-intent	-0.012* (0.007)	-0.023** (0.011)	-6.941** (3.206)	0.007*** (0.002)	0.007*** (0.002)	1.250*** (0.430)
Observations	78,376	78,376	78,376	135,775	135,775	135,775
Agent FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.123	0.125	0.133	0.105	0.105	0.103
Control Mean	0.036	0.050	12.991	0.019	0.019	3.991

Panel B: Sensitivity of Sales to Predicted Score

	Active Tab			Silent Tab		
	Any Policy (1)	# Policy (2)	Tot. Premium (3)	Any Policy (4)	# Policy (5)	Tot. Premium (6)
Treat × Score	0.001** (0.001)	0.002** (0.001)	0.556* (0.322)	-0.000 (0.006)	-0.000 (0.006)	0.320 (1.380)
Score	-0.000 (0.001)	-0.001 (0.001)	-0.162 (0.229)	0.027*** (0.005)	0.027*** (0.005)	4.951*** (1.045)
Observations	78,376	78,376	78,376	135,775	135,775	135,775
Agent FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.123	0.125	0.133	0.104	0.104	0.102
Control Mean	0.036	0.050	12.991	0.019	0.019	3.991

Panel C: Sensitivity of Sales to Screen Rank

	Active Tab			Silent Tab		
	Any Policy (1)	# Policy (2)	Tot. Premium (3)	Any Policy (4)	# Policy (5)	Tot. Premium (6)
Treat × Rank	0.0001 (0.0001)	0.0001 (0.0001)	0.0151 (0.0278)	0.0001** (0.0000)	0.0001** (0.0000)	0.0086 (0.0085)
Rank	-0.0003*** (0.0000)	-0.0004*** (0.0001)	-0.1176*** (0.0194)	-0.0002*** (0.0000)	-0.0002*** (0.0000)	-0.0305*** (0.0060)
Observations	78,376	78,376	78,376	135,775	135,775	135,775
Agent FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.124	0.126	0.134	0.104	0.104	0.102
Control Mean	0.036	0.050	12.991	0.019	0.019	3.991

**Note:** This table tests whether information provision changes the sensitivity of sales performance to visitors' predicted purchase intent by comparing the baseline sensitivity for the control group (coefficients on *High-intent* and *Middle-intent*) with the increase in sensitivity for the treatment group (coefficients on *Treat × High-intent* and *Treat × Middle-intent*). The coefficient on *Treat* represents the effect on sensitivity change for the omitted *Low-intent* group. The predictive algorithm displays two information components on the app: the high-, middle-, low-intent tags based on raw predicted score (Panel A) and a continuous % score computed as the 7-day-on-7-day change of raw predicted score representing the trend of purchase intent (Panel B). Visitors with a positive (negative) % score are displayed under the Active (Silent) Tab. Screen rank (Panel C) is not an output of the algorithm but a feature of the screen display where visitors are displayed in descending order of the continuous % score. The unit of observation is an agent-visitor. For each visitor, dependent variables are measured from the day the agent sees the visitor for the first time (i.e., when the agent enters the app page and the visitor is displayed on the screen on that day) until November 16 (or a day earlier if the agent leaves the platform before November 16). Dependent variables (except for *Any Policy*) are winsorized at the 1st and 99th percentiles. Independent variables *High-intent*, *Middle-intent*, *Score*, and *Rank* are the visitor's baseline information seen by the agent on the first day. *Score* is displayed on the app screen as a % and divided by 100 in regressions for presenting coefficients. *Rank* is capped at 100. All models include baseline control variables at the agent-visitor level, including total number of policies, total premium up until the day before the agent entered the app page and saw the visitor, and the number of days of the observational window for each agent-visitor. All models include agent fixed effects. Standard errors are clustered at the agent level. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.

**Table A.3: Information Treatment Effects on Sensitivity of Sales to Visitors' Predicted Purchase Intent: 14-day Active and Silent Tabs**

Panel A: Sensitivity of Sales to Predicted Intent Tag

	Active Tab			Silent Tab		
	Any Policy (1)	# Policy (2)	Tot. Premium (3)	Any Policy (4)	# Policy (5)	Tot. Premium (6)
Treat × High-intent	0.005 (0.008)	0.011 (0.011)	1.501 (3.515)	-0.002 (0.003)	-0.002 (0.003)	0.016 (0.321)
Treat × Middle-intent	0.003 (0.008)	0.007 (0.011)	0.647 (3.487)	-0.001 (0.002)	-0.001 (0.002)	-0.106 (0.223)
High-intent	0.014** (0.006)	0.015 (0.009)	4.382 (2.786)	0.013*** (0.003)	0.013*** (0.003)	1.291*** (0.251)
Middle-intent	0.005 (0.006)	0.006 (0.009)	0.679 (2.780)	0.006*** (0.001)	0.006*** (0.001)	0.672*** (0.162)
Treat	-0.009 (0.007)	-0.016 (0.011)	-2.757 (3.454)	-0.001 (0.001)	-0.001 (0.001)	-0.195 (0.176)
Observations	68,612	68,612	68,612	140,926	140,926	140,926
R-squared	0.013	0.014	0.015	0.020	0.020	0.021
Control Mean	0.036	0.049	13.427	0.015	0.015	1.859

Panel B: Sensitivity of Sales to Predicted Score

	Active Tab			Silent Tab		
	Any Policy (1)	# Policy (2)	Tot. Premium (3)	Any Policy (4)	# Policy (5)	Tot. Premium (6)
Treat × Score	0.001 (0.001)	0.002* (0.001)	0.480* (0.282)	-0.002 (0.005)	-0.002 (0.005)	-0.193 (0.618)
Score	0.000 (0.000)	-0.000 (0.001)	-0.301 (0.196)	0.031*** (0.004)	0.031*** (0.004)	3.515*** (0.498)
Treat	-0.006** (0.003)	-0.009** (0.004)	-2.356** (1.148)	-0.002* (0.001)	-0.002* (0.001)	-0.261** (0.132)
Observations	68,612	68,612	68,612	140,926	140,926	140,926
R-squared	0.013	0.013	0.014	0.020	0.020	0.021
Control Mean	0.036	0.049	13.427	0.015	0.015	1.859

Panel C: Sensitivity of Sales to Screen Rank

	Active Tab			Silent Tab		
	Any Policy (1)	# Policy (2)	Tot. Premium (3)	Any Policy (4)	# Policy (5)	Tot. Premium (6)
Treat × Rank	0.0000 (0.0001)	0.0000 (0.0001)	0.0092 (0.0297)	0.0000 (0.0000)	0.0000 (0.0000)	0.0011 (0.0026)
Rank	-0.0004*** (0.0001)	-0.0005*** (0.0001)	-0.1888*** (0.0213)	-0.0001*** (0.0000)	-0.0001*** (0.0000)	-0.0151*** (0.0020)
Treat	-0.0041* (0.0024)	-0.0056 (0.0035)	-1.3181 (1.1341)	-0.0022* (0.0012)	-0.0022* (0.0012)	-0.2519* (0.1451)
Observations	68,612	68,612	68,612	140,926	140,926	140,926
R-squared	0.014	0.015	0.016	0.020	0.020	0.021
Control Mean	0.036	0.049	13.427	0.015	0.015	1.859

**Note:** This table tests whether information provision changes the sensitivity of sales performance to visitors' predicted purchase intent by comparing the baseline sensitivity for the control group (coefficients on *High-intent* and *Middle-intent*) with the increase in sensitivity for the treatment group (coefficients on *Treat × High-intent* and *Treat × Middle-intent*). The coefficient on *Treat* represents the effect on sensitivity change for the omitted *Low-intent* group. The predictive algorithm displays two information components on the app: the high-, middle-, low-intent tags based on raw predicted score (Panel A) and a continuous % score computed as the 14-day-on-14-day change of raw predicted score representing the trend of purchase intent (Panel B). Visitors with a positive (negative) % score are displayed under the Active (Silent) Tab. Screen rank (Panel C) is not an output of the algorithm but a feature of the screen display where visitors are displayed in descending order of the continuous % score. The unit of observation is an agent-visitor. For each visitor, dependent variables are measured from the day the agent sees the visitor for the first time (i.e., when the agent enters the app page and the visitor is displayed on the screen on that day) until November 16 (or a day earlier if the agent leaves the platform before November 16). Dependent variables (except for *Any Policy*) are winsorized at the 1st and 99th percentiles. Independent variables *High-intent*, *Middle-intent*, *Score*, and *Rank* are the visitor's baseline information seen by the agent on the first day. *Score* is displayed on the app screen as a % and divided by 100 in regressions for presenting coefficients. *Rank* is capped at 100. All models include baseline control variables at the agent-visitor level, including total number of policies, total premium up until the day before the agent entered the app page and saw the visitor, and the number of days of the observational window for each agent-visitor. Standard errors are clustered at the agent level. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.

**Table A.4: Heterogeneity in Information Treatment Effects on Sensitivity of Sales to Visitors' Predicted Purchase Intent: By Agents' Information Processing Capacity**

Panel A: By Education

	Active Tab			Silent Tab		
	Any Policy (1)	# Policy (2)	Tot. Premium (3)	Any Policy (4)	# Policy (5)	Tot. Premium (6)
Treat × High-intent × College	-0.021 (0.017)	-0.031 (0.026)	-8.313 (7.382)	0.007 (0.007)	0.007 (0.007)	2.257 (1.580)
Treat × Middle-intent × College	-0.033** (0.015)	-0.048** (0.024)	-12.330* (6.781)	0.003 (0.006)	0.003 (0.006)	0.929 (1.237)
Treat × High-intent	0.018 (0.011)	0.031** (0.016)	7.973* (4.709)	-0.002 (0.004)	-0.002 (0.004)	-0.356 (0.904)
Treat × Middle-intent	0.017* (0.010)	0.039** (0.015)	7.268 (4.542)	-0.002 (0.004)	-0.002 (0.004)	-0.442 (0.803)
High-intent × College	0.023 (0.015)	0.029 (0.024)	11.028* (6.410)	0.003 (0.005)	0.003 (0.005)	0.488 (1.106)
Middle-intent × College	0.030** (0.013)	0.039* (0.021)	12.517** (5.863)	0.002 (0.004)	0.002 (0.004)	0.161 (0.867)
High-intent	-0.022** (0.010)	-0.039*** (0.014)	-11.843*** (4.272)	0.014*** (0.003)	0.014*** (0.003)	2.606*** (0.680)
Middle-intent	-0.024*** (0.009)	-0.038*** (0.013)	-11.854*** (4.078)	0.005** (0.003)	0.005** (0.003)	0.873 (0.573)
Observations	78,376	78,376	78,376	135,775	135,775	135,775
Agent FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.124	0.126	0.134	0.105	0.105	0.103
Control Mean	0.036	0.050	12.991	0.019	0.019	3.991

Panel B: By Work Experience

	Active Tab			Silent Tab		
	Any Policy (1)	# Policy (2)	Tot. Premium (3)	Any Policy (4)	# Policy (5)	Tot. Premium (6)
Treat × High-intent × Experienced	-0.033* (0.018)	-0.057** (0.026)	-4.096 (7.943)	-0.002 (0.007)	-0.002 (0.007)	-0.593 (1.494)
Treat × Middle-intent × Experienced	-0.024 (0.017)	-0.046* (0.024)	-5.913 (7.569)	0.006 (0.006)	0.006 (0.006)	0.961 (1.256)
Treat × High-intent	0.026* (0.014)	0.051** (0.021)	6.694 (6.510)	0.000 (0.005)	0.000 (0.005)	0.597 (1.032)
Treat × Middle-intent	0.018 (0.015)	0.038* (0.021)	5.621 (6.290)	-0.004 (0.004)	-0.004 (0.004)	-0.664 (0.906)
High-intent × Experienced	0.038** (0.015)	0.056** (0.023)	8.340 (6.460)	0.010** (0.005)	0.010** (0.005)	2.546** (1.053)
Middle-intent × Experienced	0.025* (0.015)	0.038* (0.022)	7.182 (6.056)	0.003 (0.004)	0.003 (0.004)	0.776 (0.888)
High-intent	-0.032*** (0.012)	-0.056*** (0.020)	-12.176** (5.020)	0.010*** (0.003)	0.010*** (0.003)	1.545** (0.720)
Middle-intent	-0.027** (0.012)	-0.044** (0.019)	-11.245** (4.632)	0.004 (0.003)	0.004 (0.003)	0.531 (0.600)
Observations	78,376	78,376	78,376	135,775	135,775	135,775
Agent FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.124	0.126	0.134	0.105	0.105	0.103
Control Mean	0.036	0.050	12.991	0.019	0.019	3.991

**Table A.4: Heterogeneity in Information Treatment Effects on Sensitivity of Sales to Visitors' Predicted Purchase Intent: By Agents' Information Processing Capacity—Continued**

Panel C: By Pre-treatment Performance

	Active Tab			Silent Tab		
	Any Policy (1)	# Policy (2)	Tot. Premium (3)	Any Policy (4)	# Policy (5)	Tot. Premium (6)
Treat × High-intent × Top Performer	-0.044** (0.017)	-0.046* (0.026)	-12.847 (8.114)	-0.010 (0.007)	-0.010 (0.007)	-2.366 (1.485)
Treat × Middle-intent × Top Performer	-0.033** (0.016)	-0.032 (0.024)	-11.339 (7.629)	-0.013** (0.006)	-0.013** (0.006)	-2.873** (1.258)
Treat × High-intent	0.039*** (0.014)	0.051** (0.021)	13.536* (6.940)	0.007 (0.005)	0.007 (0.005)	1.918* (1.113)
Treat × Middle-intent	0.027** (0.013)	0.035* (0.020)	10.575 (6.532)	0.007* (0.004)	0.007* (0.004)	1.781* (0.982)
High-intent × Top Performer	0.036** (0.015)	0.037 (0.024)	10.868 (6.987)	0.010** (0.005)	0.010** (0.005)	2.140** (1.050)
Middle-intent × Top Performer	0.024* (0.015)	0.022 (0.022)	8.344 (6.616)	0.009** (0.004)	0.009** (0.004)	1.933** (0.881)
High-intent	-0.037*** (0.012)	-0.052*** (0.020)	-15.005** (5.945)	0.009** (0.003)	0.009** (0.003)	1.411* (0.794)
Middle-intent	-0.030** (0.012)	-0.039** (0.018)	-13.060** (5.680)	0.000 (0.003)	0.000 (0.003)	-0.306 (0.691)
Observations	78,376	78,376	78,376	135,775	135,775	135,775
Agent FE	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.124	0.126	0.134	0.105	0.105	0.103
Control Mean	0.036	0.050	12.991	0.019	0.019	3.991

**Note:** This table tests heterogeneity in the information treatment effects on the sensitivity of sales to visitors' predicted purchase intent by agents' information processing capacity proxied by level of education (Panel A), work experience (Panel B), and pre-treatment sales performance (Panel C). *College* is an indicator variable that takes the value of one if an agent has a bachelor degree or beyond, and zero otherwise. *Experienced* is an indicator variable that takes the value of one if an agent has above-median work experience measured by the number of months since joining the platform when the agent entered the app page for the first time during the experiment period, and zero otherwise. *Top Performer* is an indicator variable that takes the value of one if an agent's pre-treatment sales performance (i.e., policy count) is in the top quartile among all agents, and zero otherwise. The unit of observation is an agent-visitor. For each visitor, the dependent variables are measured from the day the agent sees the visitor for the first time (i.e., when the agent enters the app page and the visitor is displayed on the screen on that day) until November 16 (or a day earlier if the agent leaves the platform before November 16). Dependent variables (except for *Any Policy*) are winsorized at the 1st and 99th percentiles. Independent variables *High-intent* and *Middle-intent* are the visitor's baseline information seen by the agent on the first day. All models include baseline control variables at the agent-visitor level, including total number of policies, total premium up until the day before the agent entered the app page and saw the visitor, and the number of days of the observational window for each agent-visitor. I also control for a visitor's screen rank. All models include agent fixed effects. Standard errors are clustered at the agent level. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.

Table A.5: **Information Treatment Effects on Visitor Composition**

	# New Visitors (1)	# Former Visitors (2)	% New Visitors (3)	# Zombie Visitors (4)
Treat	1.243*** (0.368)	0.277*** (0.095)	0.014* (0.007)	3.714** (1.761)
Observations	11,125	11,125	11,125	11,125
R-squared	0.142	0.252	0.041	0.332
Control Mean	9.131	2.868	0.500	49.397

**Note:** This table reports the information treatment effects on visitor composition, using Equation (1). The sample includes 5,430 agents in the treatment group and 5,695 agents in the control group. New visitors visit the agent for the first time after entry (i.e., when the agent enters the app page and receives treatment). Former visitors visit the agent both before and after entry. Zombie visitors visit the agent only before entry but never again after. Dependent variables (except for % *New Visitors*) are winsorized at the 1st and 99th percentiles. All models include baseline control variables, including total number of policies, total premium up until the day before agents entered the app page, and the number of days of the observational window. Robust standard errors are used. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.

Table A.6: **Information Treatment Effects on Sales Performance: By Relationship Strength**

	New Clients			Former Clients		
	# Policy (1)	Tot. Premium (2)	Avg. Premium (3)	# Policy (4)	Tot. Premium (5)	Avg. Premium (6)
Treat	0.267** (0.109)	168.454* (93.228)	87.282 (90.971)	0.023 (0.035)	58.198* (32.158)	178.665 (127.476)
Observations	11,125	11,125	4,312	11,125	11,125	1,622
R-squared	0.147	0.108	0.044	0.259	0.239	0.126
Control Mean	2.268	1339.944	983.880	0.544	433.381	1466.035

**Note:** This table reports the information treatment effects on agents' sales performance by relationship strength, using Equation (1). The sample includes 5,430 agents in the treatment group and 5,695 agents in the control group. Performance variables are split by whether the policyholder is a new or a former client. New vs. former is defined relative to when the agent enters the app page and receives treatment. *Avg. Premium* (average premium per policy) in Columns (3) and (6) are available only for agents who sold at least one policy to a visitor who satisfies the corresponding conditions during the experiment period. Dependent variables are winsorized at the 1st and 99th percentiles. All models include baseline control variables, including total number of policies, total premium up until the day before agents entered the app page, and the number of days of the observational window. Robust standard errors are used. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.



**Table A.7: Information Treatment Effects on Sensitivity of Agent Incentives to Visitors' Predicted Purchase Intent**

Dependent Variable: Selling High-commission Rate Policy

	Active Tab			Silent Tab		
	(1)	(2)	(3)	(4)	(5)	(6)
Treat × High-intent	0.011** (0.005)			0.001 (0.002)		
Treat × Middle-intent	0.006 (0.005)			0.001 (0.002)		
High-intent	0.006 (0.004)			0.009*** (0.002)		
Middle-intent	0.000 (0.004)			0.002* (0.001)		
Treat × Score		0.001** (0.000)			0.001 (0.004)	
Score		0.000 (0.000)			0.017*** (0.003)	
Treat × Rank			0.000 (0.000)			0.000 (0.000)
Rank			-0.000*** (0.000)			-0.000*** (0.000)
Treat	-0.012** (0.005)	-0.006*** (0.002)	-0.003 (0.002)	-0.003** (0.001)	-0.002* (0.001)	-0.002* (0.001)
Observations	79,563	79,563	79,563	136,861	136,861	136,861
R-squared	0.020	0.019	0.023	0.022	0.022	0.022
Control Mean	0.022	0.022	0.022	0.012	0.012	0.012

**Note:** This table tests whether information provision changes the sensitivity of agent incentives to visitors' predicted purchase intent by comparing the baseline sensitivity for the control group (coefficients on *High-intent* and *Middle-intent*) with the increase in sensitivity for the treatment group (coefficients on *Treat × High-intent* and *Treat × Middle-intent*). The coefficient on *Treat* represents the effect on sensitivity change for the omitted *Low-intent* group. The unit of observation is an agent-visitor. The dependent variable is an indicator variable that takes the value of one if the average commission rate of all policies sold to the visitor is above the sample median across all agent-visitors, and zero otherwise. Independent variables *High-intent*, *Middle-intent*, *Score*, and *Rank* are the visitor's baseline information seen by the agent on the first day. *Score* is displayed on the app screen as a % and divided by 100 in regressions for presenting coefficients. *Rank* is capped at 100. All models include baseline control variables at the agent-visitor level, including total number of policies, total premium up until the day before the agent entered the app page and saw the visitor, and the number of days of the observational window for each agent-visitor. Standard errors are clustered at the agent level. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.

**Table A.8: Information Treatment Effects on Product Variety**

	Unique # Products (1)	Product HHI (2)
Treat	0.127* (0.067)	-0.022** (0.009)
Observations	4,686	4,686
R-squared	0.267	0.120
Control Mean	2.647	0.686

**Note:** This table reports the information treatment effects on product variety, using Equation (1). The dependent variables are the unique number of products and the Herfindahl-Hirschman Index (HHI) of the number of policies sold for each product. The sample includes 4,686 agents who have sold at least one policy during the experiment period where the product information is available in the database. All models include baseline control variables, including total number of policies, total premium up until the day before agents entered the app page, and the number of days of the observational window. Robust standard errors are used. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.

**Table A.9: Information Treatment Effects on Policy Cancellation**

	# Canceled Policy (1)	Cancellation Ratio (2)
Treat	0.0094 (0.0082)	0.0006 (0.0027)
Observations	11,125	4,687
R-squared	0.020	0.000
Control Mean	0.046	0.016

**Note:** This table reports the information treatment effects on policy cancellation, using Equation (1). Dependent variables are the number of canceled policies and the cancellation ratio. In Column (1), the sample includes 5,430 agents in the treatment group and 5,695 agents in the control group. In Column (2), the sample includes 4,687 agents who sold at least one policy during the experiment period, including 2,321 agents in the treatment group and 2,366 agents in the control group. All models include baseline control variables, including total number of policies, total premium up until the day before agents entered the app page, and the number of days of the observational window. Robust standard errors are used. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.

Table A.10: **Information Treatment Effects on Selection: Using Claim Dummy**

Panel A: Without Insurer and Agent FE

	DV: Claim Dummy				
	Control-All (1)	Treat-All (2)	Treat-High (3)	Treat-Middle (4)	Treat-Low (5)
Log Insurance Amount	-0.005 (0.007)	0.019** (0.008)	0.029** (0.012)	0.025*** (0.009)	-0.021 (0.030)
Observations	3,186	2,832	1,065	1,546	214
Premium Rate	Yes	Yes	Yes	Yes	Yes
Age/Gender/Location	Yes	Yes	Yes	Yes	Yes
Insurance Type FE	Yes	Yes	Yes	Yes	Yes
Applicant-Insurant Relation FE	Yes	Yes	Yes	Yes	Yes
Insurer FE	No	No	No	No	No
Agent FE	No	No	No	No	No
R-squared	0.030	0.035	0.053	0.049	0.133
Outcome Mean	0.047	0.048	0.038	0.053	0.065

Panel B: Including Insurer FE

	DV: Claim Dummy				
	Control-All (1)	Treat-All (2)	Treat-High (3)	Treat-Middle (4)	Treat-Low (5)
Log Insurance Amount	-0.005 (0.010)	0.042*** (0.016)	0.051** (0.021)	0.060** (0.026)	-0.013 (0.029)
Observations	3,184	2,831	1,064	1,546	211
Premium Rate	Yes	Yes	Yes	Yes	Yes
Age/Gender/Location	Yes	Yes	Yes	Yes	Yes
Insurance Type FE	Yes	Yes	Yes	Yes	Yes
Applicant-Insurant Relation FE	Yes	Yes	Yes	Yes	Yes
Insurer FE	Yes	Yes	Yes	Yes	Yes
Agent FE	No	No	No	No	No
R-squared	0.045	0.059	0.100	0.091	0.151
Outcome Mean	0.047	0.048	0.038	0.053	0.062

Panel C: Including Insurer and Agent FE

	DV: Claim Dummy				
	Control-All (1)	Treat-All (2)	Treat-High (3)	Treat-Middle (4)	Treat-Low (5)
Log Insurance Amount	-0.007 (0.017)	0.055** (0.021)	0.061** (0.030)	0.066* (0.038)	-0.010 (0.036)
Observations	2,969	2,625	923	1,396	185
Premium Rate	Yes	Yes	Yes	Yes	Yes
Age/Gender/Location	Yes	Yes	Yes	Yes	Yes
Insurance Type FE	Yes	Yes	Yes	Yes	Yes
Applicant-Insurant Relation FE	Yes	Yes	Yes	Yes	Yes
Insurer FE	Yes	Yes	Yes	Yes	Yes
Agent FE	Yes	Yes	Yes	Yes	Yes
R-squared	0.209	0.265	0.230	0.279	0.592
Outcome Mean	0.045	0.047	0.039	0.050	0.054

**Note:** This table examines selection in different agent-visitor cohorts using the risk-coverage correlation model. The analysis is at the policy level. Risk is measured using a dummy variable which takes the value of one if there is any claim for a policy, and zero otherwise. Coverage is measured as the logarithm of insurance amount. The sample includes all policies sold to visitors during the experiment period where all variables used in regressions are available. *Treat-High*, *Treat-Middle*, and *Treat-Low* are based on a visitor's baseline purchase intent classification seen by the agent on the first day. Standard errors are clustered at the agent level. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.

Table A.11: AI Demand Estimates and Claim Outcomes

	DV: Log Claim Amount		
	(1)	(2)	(3)
Middle-intent	-0.087 (0.098)		
High-intent	-0.093 (0.096)		
Score		-0.001 (0.011)	
Rank			0.001 (0.001)
Observations	5,604	5,604	5,604
Log Insurance Amount	Yes	Yes	Yes
Premium Rate	Yes	Yes	Yes
Age/Gender/Location	Yes	Yes	Yes
Insurance Type FE	Yes	Yes	Yes
Applicant-Insurant Relation FE	Yes	Yes	Yes
Insurer FE	Yes	Yes	Yes
Agent FE	Yes	Yes	Yes
R-squared	0.219	0.218	0.219
Outcome Mean	0.263	0.263	0.263

**Note:** This table examines the correlation between AI demand estimates and ex-post claim outcomes. The analysis is at the policy level. The dependent variable is the logarithm of claimed amount. The sample includes all policies sold to visitors during the experiment period where all variables used in regressions are available. *High-intent*, *Middle-intent* (*Low-intent* is omitted), *Score*, and *Rank* are based on a visitor's baseline purchase intent information seen by the agent on the first day. Standard errors are clustered at the agent level. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.

Table A.12: **Information Treatment Effects on Product Selection/Indirect Pricing**

	DV: Log Premium		
	(1)	(2)	(3)
Log Claim Amount	0.016** (0.006)	0.008** (0.004)	0.008** (0.004)
Log Claim Amount × Treat	0.018 (0.015)	0.016 (0.012)	0.005 (0.006)
Treat	0.016 (0.028)	0.009 (0.015)	
Observations	6,026	6,024	5,604
Log Insurance Amount	Yes	Yes	Yes
Age/Gender/Location	Yes	Yes	Yes
Insurance Type FE	Yes	Yes	Yes
Applicant-Insurant Relation FE	Yes	Yes	Yes
Insurer FE	No	Yes	Yes
Agent FE	No	No	Yes
R-squared	0.641	0.809	0.912
Outcome Mean	5.021	5.020	5.005

**Note:** This table examines the treatment effects on agents' product selection as indirect pricing, by testing whether AI demand predictions change the risk-price correlation. The analysis is at the policy level. Risk is measured as the logarithm of claimed amount. Price is measured as the logarithm of premium. Standard errors are clustered at the agent level. \*, \*\*, and \*\*\* denote significance at the 10%, 5%, and 1% levels, respectively.