AI Adoption and the Cost of Access in Digital Health Platforms*

Jian Ni[1], Jiang Qian[2], and Meng Zhu[3]

[1]Virginia Tech
[2]University of Sydney
[3]Virginia Tech

November 2, 2025

Abstract

Although artificial intelligence (AI) has transformed the way digital platforms match users with products and services, its broader market and societal consequences remain poorly understood. We examine how AI-driven (matching) tools in teletherapy platforms — designed to connect patients with therapists more efficiently — affect pricing and healthcare access. Using data from a leading US teletherapy platform that covers approximately 150,000 therapists in all states, we exploit the temporal and cross-state variation in AI regulatory environments to identify the causal impact of AI adoption. Our analysis shows that the integration of AI tools leads to significantly higher therapy prices in both in-person and online sessions. The evidence indicates that improved matching efficiency increases the visibility and demand of the therapists without corresponding improvements in service quality, thus changing the market equilibrium. To further explore distributional effects, we employ causal forest estimation to uncover heterogeneous treatment effects across markets and population segments and find that price increases are most pronounced in counties with lower household income and fewer local providers. These results suggest that the adoption of AI improves market efficiency, but can also exacerbate affordability gaps in mental health care. Our findings highlight how AI tools that optimize individual choice can produce unintended aggregate consequences, such as increasing prices and deepening access disparities. Our study emphasizes that as AI transforms digital marketplaces, attention must be paid to how efficiency improvements reshape access, affordability, and overall welfare.

Keywords: AI Adoption; Teletherapy; Pricing; Market Access; Causal Forest

*Ni: Pamplin College of Business, Virginia Tech, VA 24060. Email: jiann@vt.edu. Qian: University of Sydney Business School, Camperdown, NSW 2006, Australia. Email: jiang.qian@sydney.edu.au. Zhu: Pamplin College of Business, Virginia Tech, VA 24060. Email: mengz@vt.edu. The usual disclaimer applies.

1 Introduction

Rapid integration of artificial intelligence (AI) into digital platforms has fundamentally reshaped how consumers and service providers connect across numerous markets. In healthcare, where information asymmetries and matching frictions have long posed significant challenges, AI-powered matching algorithms promise to revolutionize service delivery by connecting patients with providers who better fit their needs. However, while the potential benefits of improved matching are widely touted, the economic consequences of AI adoption remain poorly understood — particularly when AI enhances matching efficiency without necessarily improving underlying service quality.

Our study examine a critical yet underexplored dimension of AI adoption in healthcare: how AI-driven algorithms or tools affect pricing in online mental health therapy platforms. Unlike traditional technological innovations that improve service quality or reduce costs, AI algorithms or tools operate through a different mechanism. By more efficiently connecting patients with therapists whose characteristics align with patient preferences, these tools increase the effective demand for individual therapists without altering the therapeutic services they provide. This demand-driven price effect presents a puzzle: can technological improvements that enhance market efficiency paradoxically increase costs and exacerbate access disparities?

The mental health therapy market provides an ideal setting to investigate this phenomenon. Telehealth platforms have experienced explosive growth, particularly following the COVID-19 pandemic, with mental health services representing one of the largest segments of the telemedicine market. These platforms increasingly deploy sophisticated AI algorithms to match patients with therapists based on specialty, availability, communication style, and demographic characteristics. Therapists, in turn, may adjust their pricing in response to increased demand generated by improved algorithmic matching. Crucially, this price response occurs independently of any changes in service quality—the therapeutic hour remains substantively unchanged, yet its market value shifts due to algorithmic mediation.

Our research addresses two fundamental questions: First, does AI adoption in matching algorithms causally affect therapy pricing? Second, are these pricing effects heterogeneous across geographic areas with different socioeconomic characteristics? Answering these questions is empirically challenging. The decision to adopt AI-driven matching is endogenous — platforms and therapists may adopt AI in response to market conditions that also influence pricing. Moreover, unobserved factors such as platform quality, regional demand trends, and therapist characteristics could confound the relationship between AI adoption and prices.

To address these identification challenges, we employ a Spatial Regression Discontinuity

Design that exploits the staggered adoption of state-level AI legislation as a natural experiment. Focusing on border counties across states with differential AI regulatory environments, we isolate exogenous variation in AI adoption while holding constant many unobserved regional factors. We analyze a comprehensive dataset of nearly 150,000 therapist profiles from a major U.S. teletherapy platform, implementing a Two-Way Fixed Effects framework to estimate causal effects. Additionally, we deploy causal forest methods to uncover heterogeneous treatment effects across counties with varying socioeconomic profiles.

Our findings reveal that AI adoption significantly increases therapy prices, with effects concentrated in economically vulnerable areas. Specifically, counties with lower median incomes and lower educational attainment experience substantially larger price increases following AI adoption. This pattern suggests that while AI matching algorithms improve market efficiency by reducing search frictions, they simultaneously generate distributional consequences that may deepen healthcare access disparities. The mechanism appears clear: better matching increases demand for individual therapists without expanding supply or improving quality, creating upward pressure on prices that disproportionately affects cost-sensitive populations.

These results contribute to several streams of literature. First, we extend research on AI adoption in healthcare by documenting a novel channel through which AI affects market outcomes—not through quality improvements or cost reductions, but through demand reallocation. Second, we contribute to the growing literature on algorithmic pricing and platform economics by demonstrating how matching algorithms can indirectly influence prices through demand effects. Third, we add to the telemedicine literature by highlighting an unintended consequence of technological innovation: efficiency improvements may inadvertently create new barriers to access.

The policy implications are significant. As AI-driven matching becomes ubiquitous in digital health platforms, policymakers and platform designers must consider not only the efficiency gains from improved matching but also the distributional consequences of demand-driven price inflation. Our findings suggest that without deliberate intervention—such as price controls, subsidies for low-income populations, or alternative platform designs—AI adoption may paradoxically worsen healthcare equity even as it improves matching efficiency.

The remainder of this paper proceeds as follows. Section 2 reviews relevant literature on telemedicine, AI adoption, and spatial identification strategies. Section 3 provides institutional context on the mental health therapy market and AI integration in telehealth platforms. Section 4 describes our data from a leading U.S. teletherapy platform. Section 5 outlines our empirical methodology, including the spatial regression discontinuity design

and causal forest analysis. Section 6 presents our main findings and heterogeneity analysis. Section 7 discusses implications and concludes.

2 Literature Review

Our paper is related to the literature of AI and healthcare. Zhu, Chakravarti, and Ni (2022) examine the transformative role of artificial intelligence (AI) and other digital tools in transitioning healthcare toward a patient-centered and data-driven model. The authors emphasize AI's potential to enhance personalized care, improve decision-making, and address structural barriers in healthcare, while also highlighting critical challenges such as algorithmic bias, privacy concerns, and inclusivity. Their work underscores the importance of leveraging AI to balance efficiency and equity in healthcare delivery. Moorman, Van Heerde, Moreau, and Palmatier (2024) argue that AI can enhance healthcare access and efficiency by driving productivity improvements and fostering competition. AI-enabled technologies, such as wearables and diagnostic tools, are increasingly adopted by healthcare providers, which could theoretically reduce service costs while improving accessibility.

Despite these theoretical benefits, the practical implementation of AI in healthcare is complex. Providers may resist adopting AI tools, even when they enhance efficiency, due to reputational or ethical concerns (Karimian, Petelos, and Evers, 2022). Furthermore, consumers' perceptions of AI adoption in healthcare can significantly influence providers' pricing strategies. While AI's promise to enhance efficiency and expand access is well-documented, the actual cost implications remain underexplored. This underscores the importance of investigating the effects of AI adoption on pricing, particularly in specialized areas like mental health services where disparities in access and affordability are already prominent.

Our paper is also related the literature of border strategy which leverages geographic discontinuities to estimate causal effects by comparing outcomes on either side of a boundary where policies, advertisements, or other factors differ. Holmes (1998) used border discontinuities to evaluate the economic impacts of right-to-work laws on manufacturing activity, showcasing how policy variations at borders create natural experiments for causal inference. In marketing, border strategies are frequently employed to assess the effectiveness of advertising campaigns. Wang, Lewis, and Schweidel (2018) used a border strategy to examine the differential impacts of political advertisements aired across designated marketing areas (DMAs). Shapiro (2016) employed a similar approach in pharmaceutical marketing to evaluate positive spillover effects of antidepressant advertising. This study highlighted how discontinuities in television-market borders allowed for the estimation of advertising's influ-

ence on both own and rival demand, addressing endogeneity concerns by leveraging spatial variation.

The border strategy is particularly effective in isolating localized treatment effects, providing robust evidence in settings where traditional experimental designs are infeasible. Our study employs such a border strategy similar to the above papers to estimate the causal effects of AI adoption in telemedicine. By leveraging geographic discontinuities created by state-level AI legislation, we compare therapy pricing in border counties where regulations differ, isolating the impact of AI adoption on service costs. This approach allows us to address endogeneity concerns and account for unobserved heterogeneity, offering a robust framework for evaluating the economic and equity implications of AI integration in healthcare.

Our paper contributes to the growing body of research examining telemedicine which started to transform healthcare delivery, especially in mental health services, by expanding access and increasing patient choice. Goetz (2023) examines the impact of telemedicine competition on incumbent healthcare providers in the Canadian psychotherapy market. Using a difference-in-differences approach, the study finds that new telemedicine entrants caused incumbents to stop offering income-based discounts but did not significantly alter their base prices. This suggests that telemedicine competition may increase costs for price-sensitive patients while benefiting those with higher ability to pay. The findings highlight the role of capacity constraints and price discrimination strategies in telemedicine markets. Zhou et al. (2021) explore the implications of telemedicine for rural healthcare systems, showing that urban telemedicine providers gained market share at the expense of rural hospitals. This highlights the disruptive potential of telemedicine in redistributing patient volumes and revenue, particularly in underserved areas.

As telemedicine platforms evolve, the integration of AI has emerged as a key factor in enhancing efficiency and decision-making. AI-driven tools enable personalized patient-provider matching and streamline administrative processes, potentially improving patient satisfaction and reducing operational costs. However, the adoption of AI in telemedicine is not without challenges. Providers may resist using AI tools due to reputational concerns, privacy issues, and biases in algorithmic decision-making (Tóth et al., 2022).

The intersection of telemedicine and AI adoption remains underexplored in terms of its economic and equity implications. While AI promises to improve accessibility and affordability, its integration may inadvertently exacerbate disparities, particularly in low-income regions. Studies like Goetz (2023) underscore the need for nuanced policies that balance the benefits of AI-driven telemedicine with the risks of inequality and market disruption. Our study contributes to the telemedicine literature by providing the first causal evidence on

how AI adoption impacts therapy pricing and healthcare equity. While existing research has primarily focused on telemedicine's role in expanding access and its competitive effects on incumbent providers, our work highlights the nuanced economic and social consequences of integrating AI tools into telemedicine platforms. By leveraging a border strategy and exploring heterogeneous impacts across socioeconomic groups, we bridge gaps in understanding the interplay between technological innovation and disparities in healthcare affordability.

3 Background

3.1 Prevalence of Mental Health Disorders

Mental health disorders constitute a significant public health concern in the US. In the year 2021, approximately 57.8 million adults, aged 18 or older, were reported to have experienced some form of mental illness in the United States. This figure accounted for 22.8% of the total adult population in the U.S. (Ringeisen et al., 2023). Moreover, the prevalence of mental health disorders among adolescents and children is also prevalent. Anxiety disorders were the most prevalent among U.S. adolescents (31.9%), followed by behavior disorders (19.1%), mood disorders (14.3%), and substance use disorders (11.4%). Overall, one in four to five youths in the U.S. experiences a mental disorder with severe impairment throughout their lifetime (Merikangas et al., 2010).

3.2 Access to Mental Health Services

While the prevalence of mental health disorders has garnered substantial attention in recent years due to its critical importance in public health, Access to mental health services remains a challenge. A substantial treatment gap persists, with many individuals, particularly those from marginalized communities, facing barriers in obtaining adequate care (Alegría et al., 2021). Factors such as stigma, limited availability of providers, and socioeconomic disparities contribute to this access gap.

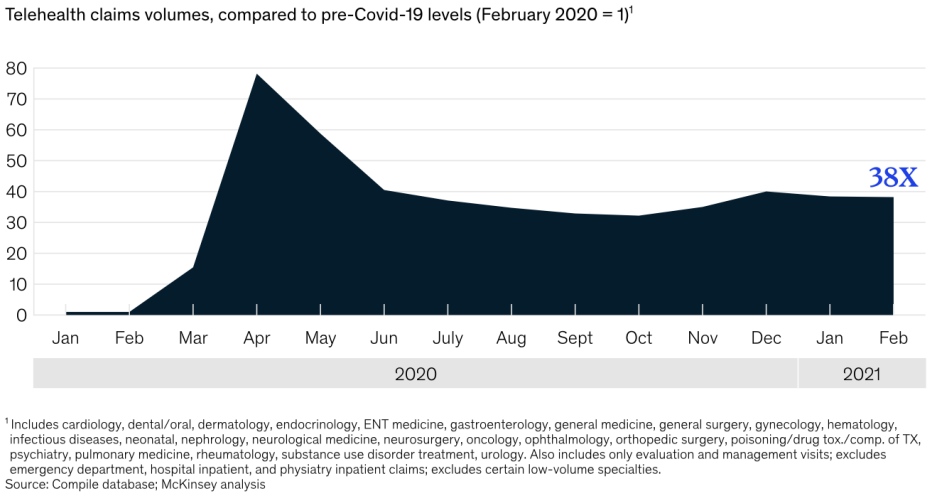
3.3 Role of Telehealth in Mental Health Care

Advancements in technology have played a significant role in the mental health market, particularly with the introduction of teletherapy. Teletherapy utilizes a range of interactive features, including video conferencing, audio, and text-based messaging, to facilitate the

therapeutic relationship. These features provide patients with a sense of connection and engagement similar to in-person sessions. Telehealth also offers the advantage of greater flexibility in scheduling, which is particularly beneficial for individuals who struggle to take time off work or face transportation issues. Additionally, telehealth offers the benefit of anonymity and privacy, addressing the concerns of individuals who may be hesitant to seek in-person treatment.

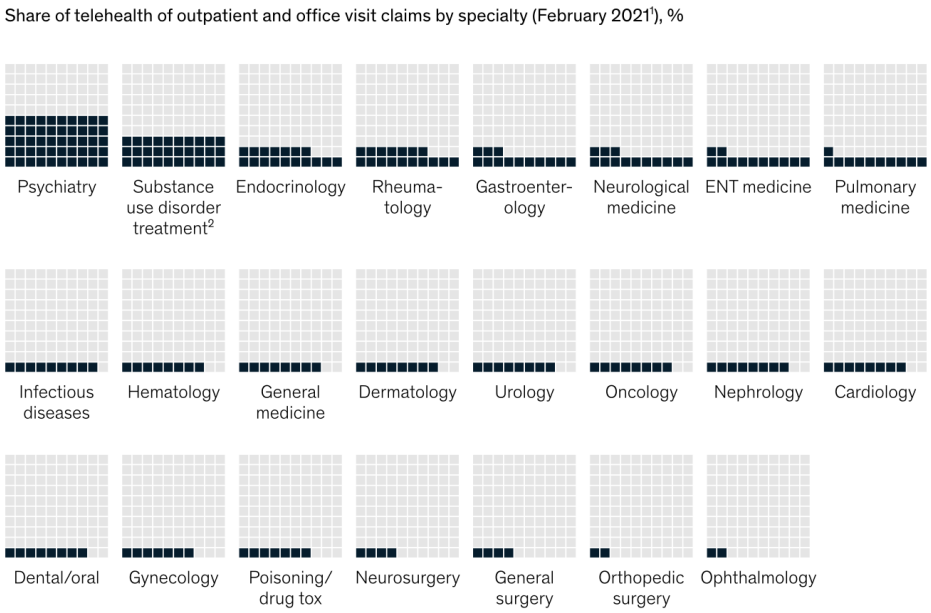
Telehealth has emerged as a crucial tool in expanding access to mental health care. The COVID-19 pandemic accelerated the adoption of telehealth, with a rapid increase in virtual mental health visits. As reported by McKinsey (Bestsennyy et al., 2021), telehealth utilization has demonstrated remarkable resilience, maintaining levels that are 38 times higher than the pre-pandemic era (See Figure 1). This sustained adoption accounts for over two-thirds of the originally projected virtualizable visits. In the wake of the initial surge witnessed in April 2020, telehealth adoption has progressively climbed to encompass up to 17 percent of all outpatient and office visit claims, particularly those involving evaluation and management services. This steady integration has been observed since June 2020. Furthermore, there exists a discernible variation in telehealth adoption across medical specialties, with psychiatry leading the way at 50 percent, closely followed by substance use disorder treatment at 30 percent (See Figure 2). This shift to telehealth facilitated broader reach, particularly for individuals in remote or underserved areas.

Figure 1: Growth in telehealth usage peaked during April 2020 but has since stabilized[a]



[a]Source: (Bestsennyy et al., 2021)

Figure 2: Substantial variation exists in share of telehealth claims across specialties[a]



3.4 The Effectiveness of Telehealth

Research has demonstrated that virtual treatment can be as effective as in-person treatment for various mental health conditions, such as depression, anxiety, and post-traumatic stress disorder ((Turgoose et al., 2018)). Furthermore, studies ((Hunter et al., 2022)) have highlighted patients who face difficulties accessing in-person care, such as those living in rural or underserved areas, older adults, or individuals with mobility limitations, are keend for virtual treatment. Moreover, Bestsennyy et al. (2021) reported that Approximately 40 percent of surveyed consumers express their intention to persist in using telehealth as part of their ongoing healthcare journey—a substantial increase from the 11 percent who utilized telehealth before the COVID-19 pandemic.

3.5 Economic Burden of Mental Health Treatment

The economic burden of mental health treatment in the US is substantial. Direct costs associated with mental health treatment, as well as indirect costs related to lost productivity, account for a significant portion of healthcare expenditures ((Roehrig, 2016)). These direct costs contribute significantly to the financial burden borne by individuals, families, and the

healthcare system as a whole. In addition to the direct treatment costs, the economic impact of mental health treatment extends to indirect costs, primarily driven by lost productivity in the workforce. Untreated or undermanaged mental health conditions can lead to reduced workplace efficiency, increased absenteeism, and diminished overall productivity ((Stewart et al., 2003)). This translates into substantial economic losses for employers and the broader economy.

3.6 Transforming Doctor-Patient Matchmaking in Telehealth

Despite the advancing realm of virtual healthcare, contemporary telehealth platforms often adhere to a conventional directory-style approach. While telehealth platforms make it easier for patients to explore more therapists, the challenge of finding the right doctors who matches their specific situation still persists. Patients navigate static web pages to discover doctors' information, employing search criteria like geographical locations or specific mental health concerns. However, this method frequently leaves patients feeling uncertain when it comes to selecting a doctor online. In response to this challenge, teletherapy platforms are turning to AI technology to empower patients in making more informed decisions.

The mental health market has witnessed significant advancements with the integration of AI tools in teletherapy platforms. By leveraging the capabilities of AI, teletherapy platforms strive to elevate the patient experience by offering dependable and tailored suggestions. AI-driven tools can take into account an array of factors, including therapy specializations, patient preferences, to present personalized therapist matches. This infusion of AI technology holds the potential to reshape the way patients connect with mental health professionals, thereby enhancing the accessibility and efficacy of therapy services within the dynamic landscape of the mental health market.

3.7 Costs of Integrating AI in Teletherapy Platform

Although platforms invest significant effort in incorporating AI for the benefit of both customers and therapists, it's crucial to note that therapists' consent remains essential. Activating the AI system involves therapists in creating detailed text and video profiles, which are then submitted for assessment. The AI algorithm utilizes this information to recommend a curated selection of therapists to patients. Once AI is activated, therapists' profiles feature a richer array of information, including text, video, and audio components, thereby attracting a larger pool of potential patients. While the activation of AI functionality comes at no cost, therapists are required to dedicate time to consistently refine and update their profile

information.

3.8 Balancing AI Potential with Relevant Policy

Although AI holds great promise across various sectors, policymakers are increasingly vigilant about safeguarding consumer privacy and ensuring responsible AI use. This vigilance has prompted a surge in state-level regulations governing AI applications. In AI-related legislations enacted in each state up until 2022, New York State takes the lead with more than five enacted measures. Other states have also introduced at least one legislation concerning AI. The regulatory landscape varies, with some states adopting cautious stances to address potential biases and discrimination, while others incentivize AI-driven enterprises through tax benefits.

One notable example is the Illinois AI Video Interviewing Act 820 ILCS 42, which mandates employers to transparently disclose the use of AI in analyzing interviews and evaluating candidates for positions. Employers must provide detailed information about the functioning of AI, the attributes it assesses, secure consent from applicants, and facilitate the sharing of interview videos. Thus, the ethical considerations and concerns regarding privacy and data security are important areas of discussion ((Luxton et al., 2014)). Given that the teletherapy platform under scrutiny necessitates personal videos from therapists to enable AI-powered recommendations, therapists' decisions to embrace AI are directly shaped by these legislations and apprehensions regarding the handling of their personal data.

3.9 Concerns Regarding AI Adoption

While AI-powered tools that connect therapists with patients are available for adoption at no cost, therapists must manually activate them after furnishing text and video-based information. Although this AI technology streamlines the process of matching patients with suitable therapists, it comes with potential drawbacks. One potential concern with AI is its potential exploitation as a marketing ploy to inflate the costs of teletherapy sessions. On one hand, it bolsters the efficiency of patient-therapist pairings, leading to heightened demand. This enables therapists to target specific patient demographics, potentially those with a higher capacity to pay. Consequently, the surge in demand may exert upward pressure on the fees charged by therapists who utilize AI tools for self-promotion, potentially exacerbating healthcare disparities.

In sum, while the aim of AI adoption is to enhance the overall patient experience, its

impact on therapy service pricing remains a relatively unexplored area. This study endeavors to bridge this research gap, delving into the effects of AI adoption on therapy service costs, offering valuable insights into this burgeoning field.

4 Data Description

Our research utilizes data from one of the leading teletherapy platforms in the United States, collected in December 2021 and June 2022. The platform had a total of 142,580 registered psychologists representing all 50 states and Washington, D.C. in December 2021, which increased to 155,686 by June 2022. This dataset provides a robust and comprehensive view of the current state of teletherapy in the United States, encompassing therapists' decisions to adopt AI, service pricing, therapists' profiles, and other relevant details. Analyzing data from different time points allows us to observe changes and trends in the market, providing valuable insights.

4.1 Therapy Price

The therapy prices in our study are categorized into virtual treatment and in-person treatment. Virtual treatment is charged at a fixed rate, such as $100 per session, while in-person treatment has a price range, with a low and high values (e.g., [80-120] represents a low price of $80 per session and a high price of $120 per session). Since the exact percentage of low and high prices in actual therapist transactions is unknown, our analysis will be conducted on each of the three price categories separately. Summary statistics for each price category are presented in Table 1.

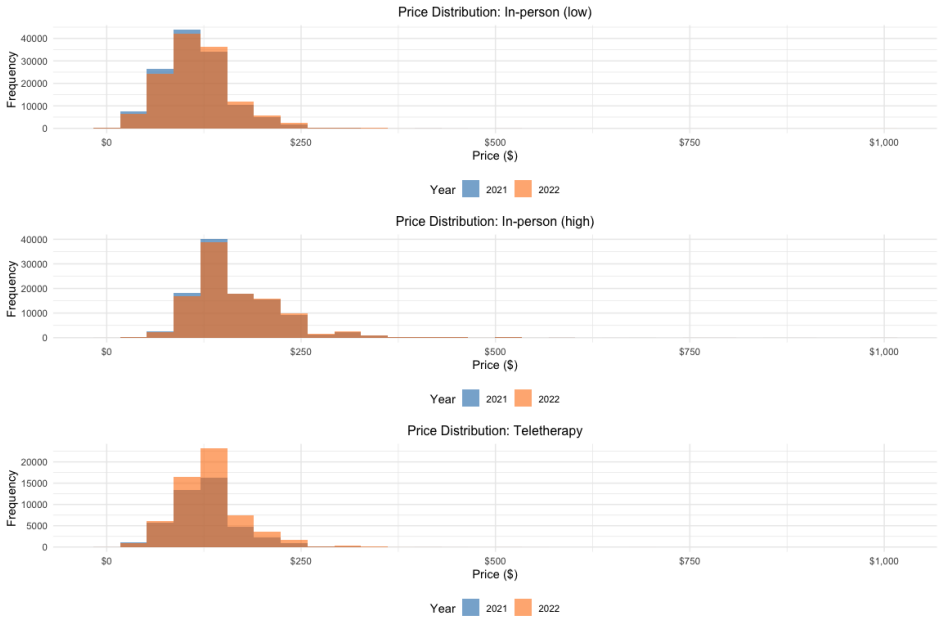
Table 1: The Summary Statistics of the Therapy Treatment Price

Date	Price	N	Mean	SD	25%	Median	75%	Min	Max
2021	in-person (low)	129,859	116.50	95.27	85	110	150	1	14,500
	in-person (high)	109,110	170.58	219.58	130	150	200	1	30,000
	teletherapy	44,791	151.80	1,112.26	100	125	150	1	150,175
2022	in-person (low)	130,184	120.45	112.50	90	120	150	1	19,500
	in-person (high)	107,728	173.55	179.06	130	150	200	1	30,000
	teletherapy	60,038	146.78	895.63	100	130	150	1	200,275

The histogram plots visualized in Figure 3 illustrate a right-skewed distribution for prices of both teletherapy and in-person treatments, with the majority falling within the $50-$200 range. Specifically, the averages stand at $149.3 for teletherapy and $118.5 for in-person

treatment when pooling data across both years. Comparing prices between 2021 and 2022 reveals mixed trends across treatment types: teletherapy prices decreased slightly by $5.02 (approximately 3.3%), while in-person treatment prices increased for both categories—low prices rose by $3.95 (approximately 3.4%) and high prices increased by $2.97 (approximately 1.7%).

Figure 3: The Histogram of the Therapy Treatment Price



Upon closer examination of the pricing data, we find that the vast majority of therapy prices fall within reasonable and expected ranges, demonstrating high data quality. While we observed some extreme values, including a minimum treatment price of $1, these cases represent only a negligible portion of the dataset. Specifically, prices of $1 or below constitute merely 0.023% of in-person (low) prices, 0.006% of in-person (high) prices, and 0.021% of teletherapy prices. More broadly, prices below $30—which might be considered unusually low for professional therapy services—account for only 1.336% of in-person (low) prices, 0.050% of in-person (high) prices, and 0.300% of teletherapy prices. At the upper end, extremely high prices above $1,000 represent less than 0.132% across all categories. This distribution indicates that over 98% of all therapy prices in our dataset fall within realistic market ranges, supporting the reliability and validity of our pricing data for subsequent analyses.

4.2 AI Adoption

The analysis of AI adoption patterns reveals significant temporal growth and substantial geographic heterogeneity in therapist technology adoption behaviors. AI adoption rates experienced measurable growth between our two observation periods, increasing from 8.71% in 2021 to 9.16% in 2022, representing a 5.2% relative increase over time (See Table 2). While this growth rate may appear modest in absolute terms, it represents meaningful momentum in the mental health sector, where technology adoption has historically been slower compared to other healthcare domains.

Table 2: AI Adoption Summary Statistics

Date	Total Therapists	AI Adopters	Non-Adopters	Adoption Rate (%)
2021	168,252	14,659	153,593	8.71%
2022	176,453	16,159	160,294	9.16%

The state-level analysis reveals striking disparities in AI adoption that underscore fundamental differences in regional attitudes toward healthcare technology integration (Table 3, Figure 4). AI adoption rates exhibit a large variation across states, ranging from a minimum of 4% in Vermont to a maximum of 12% in Colorado and Florida. This substantial heterogeneity suggests that local factors—including regulatory environments, professional culture, population demographics, and technological infrastructure—play critical roles in shaping adoption decisions. Several western and southern states demonstrate relatively higher adoption rates, with Colorado and Florida leading at 12%, followed by Texas and Missouri at 11%. Conversely, several traditionally technology-forward northeastern states show surprisingly lower adoption rates, including Vermont (4%), Massachusetts (5%), New Hampshire (5%), and Montana (5%).

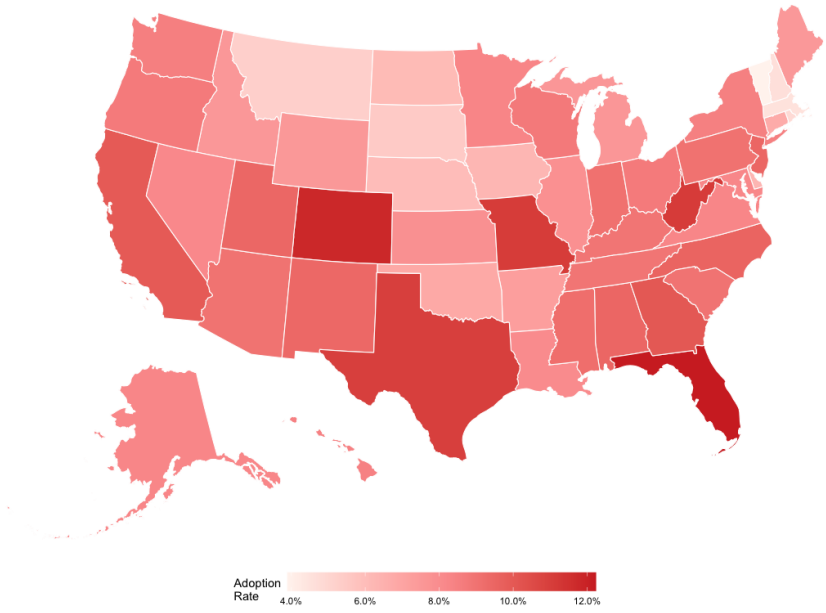
When examining the relationship between therapist density and adoption rates, the data reveals interesting market dynamics that challenge conventional technology diffusion assumptions. States with varying therapist concentrations show different adoption patterns that reflect local market characteristics. For instance, Vermont has the highest therapist density at 3,455.56 therapists per million residents but shows the lowest adoption rate at 4%, while Colorado with 2,513.00 therapists per million achieves the highest adoption rate at 12%. Similarly, West Virginia, with a relatively low therapist density of 337.50 per million, maintains a strong 11% adoption rate, while Massachusetts with 1,514.43 therapists per million shows only 5% adoption.

The geographic distribution of adoption rates suggests that AI technology serves multiple market functions rather than simply addressing resource scarcity. High-adoption states

Table 3: AI Adoption Rate by State

State	Population (million)	Adoption	Therapists	Adoption Rate	Therapists Per (million)
Alabama	4.86	221	2333	0.09	480.04
Alaska	0.74	55	665	0.08	898.65
Arizona	6.83	631	7000	0.09	1024.89
Arkansas	2.98	163	2241	0.07	752.01
California	39.14	2099	21059	0.10	538.04
Colorado	5.46	1618	13721	0.12	2513.00
Connecticut	3.59	591	8762	0.07	2440.67
Delaware	0.95	86	1357	0.06	1428.42
District of Columbia	0.67	288	3697	0.08	5517.91
Florida	20.27	2535	20729	0.12	1022.64
Georgia	10.21	1210	12040	0.10	1179.24
Hawaii	1.43	142	1679	0.08	1174.13
Idaho	1.65	173	2296	0.08	1391.52
Illinois	12.86	1174	14968	0.08	1163.92
Indiana	6.62	377	4153	0.09	627.34
Iowa	3.12	152	2408	0.06	771.79
Kansas	2.91	242	3078	0.08	1057.73
Kentucky	4.43	269	3017	0.09	681.04
Louisiana	4.67	301	3732	0.08	799.14
Maine	1.33	195	2603	0.07	1957.14
Maryland	6.01	973	11871	0.08	1975.21
Massachusetts	6.79	466	10283	0.05	1514.43
Michigan	9.92	983	12930	0.08	1303.43
Minnesota	5.49	248	2983	0.08	543.35
Mississippi	2.99	124	1351	0.09	451.84
Missouri	6.08	617	5561	0.11	914.64
Montana	1.03	102	1943	0.05	1886.41
Nebraska	1.90	131	2170	0.06	1142.11
Nevada	2.89	245	3001	0.08	1038.41
New Hampshire	1.33	96	2076	0.05	1560.90
New Jersey	8.96	1347	14214	0.09	1586.38
New Mexico	2.09	182	1935	0.09	925.84
New York	19.80	1888	22181	0.09	1120.25
North Carolina	10.04	1426	14901	0.10	1484.16
North Dakota	0.76	20	327	0.06	430.26
Ohio	11.61	724	8306	0.09	715.42
Oklahoma	3.91	196	2867	0.07	733.25
Oregon	4.03	691	7929	0.09	1967.49
Pennsylvania	12.80	1373	15155	0.09	1183.98
Rhode Island	1.06	99	2044	0.05	1928.30
South Carolina	4.90	378	4206	0.09	858.37
South Dakota	0.86	25	457	0.05	531.40
Tennessee	6.60	562	6297	0.09	954.09
Texas	27.47	2279	20774	0.11	756.24
Utah	3.00	446	4720	0.09	1573.33
Vermont	0.63	85	2177	0.04	3455.56
Virginia	8.38	986	11934	0.08	1424.11
Washington	7.17	1163	13581	0.09	1894.14
West Virginia	1.84	69	621	0.11	337.50
Wisconsin	5.77	289	3298	0.09	571.58
Wyoming	0.59	44	585	0.08	991.53

Figure 4: AI Adoption Rates by State



include both resource-rich markets like Colorado (2,513.00 therapists per million, 12% adoption) and resource-constrained markets like West Virginia (337.50 therapists per million, 11% adoption). This pattern indicates that AI adoption may be driven by different factors in different contexts: market efficiency and service enhancement in well-resourced areas, and capacity augmentation in underserved regions. States with moderate therapist densities between 500-1,500 per million show varied adoption rates from 5% to 11%, suggesting that local factors beyond simple resource availability influence technology integration decisions.

These findings indicate that AI adoption in mental health reflects a maturing technology market where adoption decisions are influenced by multiple factors including market sophistication, professional culture, regulatory environment, and strategic positioning rather than solely by resource constraints. The geographic patterns suggest opportunities for knowledge transfer and best practice sharing between high-adoption and low-adoption states, regardless of their underlying resource profiles. Understanding these varied adoption contexts is crucial for analyzing how technology integration affects market dynamics and pricing structures. This relationship between technology adoption and pricing has important implications for policymakers and healthcare administrators seeking to promote effective technology integration strategies while considering the economic accessibility and affordability of mental health services across different market conditions and resource environments.

4.3 Other Variables

In addition to therapy prices and AI adoption, we collected other variables that may impact pricing, such as the availability of sliding scale options and accepted payment methods and insurance plans. Sliding scale options allow prices to vary per therapist, while most therapists accept cash, checks, and credit cards, including popular cards like Visa and Mastercard, as well as American Express, Discover, and other virtual payment methods. The variety of payment methods is influenced by processing fees and may impact prices. Accepted insurance plans exhibit more variation compared to payment methods. We counted the number of insurance plans accepted by each therapist to exclude the potential impact of insurance on prices.

Furthermore, we collected personal demographics of therapists, including their address, years in practice, license number, graduation school, and graduation time. License numbers and graduation schools were encoded as dummy variables since the license information is a string of numbers and does not provide additional insights. Graduation school information is highly dispersed, so dummy variables were used to record whether a therapist provided relevant practice information.

To consider regional factors that could affect pricing, we obtained county-level indicators from the United States Census Bureau using therapists' addresses. These county-level factors include median household income, unemployment rate, education, and more, offering valuable insights into the socio-economic conditions of the therapists' respective areas.

5 Empirical Strategy and Results

5.1 Regression Analysis

To begin our analysis, we examine the average treatment effect of AI adoption on therapy prices using a two-way fixed effects regression model. We focus on therapists located in counties situated along state borders and investigate how state-level AI adoption influences three types of therapy prices: minimum in-person price, maximum in-person price, and teletherapy price. These prices are observed at the individual therapist level, and AI adoption is captured by a binary indicator for whether a therapist has adopted AI tools in their practice.

Let $TreatmentPrice_{ict}$ denote the price charged by therapist i in county c at time t, where t takes the value 0 for 2021 and 1 for 2022. The key independent variable, $AIAdoption_{it}$, is a binary variable equal to 1 if therapist i has adopted AI at time t, and 0 otherwise. We

estimate the following specification:

$$Ln(TreatmentPrice_{ict}) = \beta_0 + \beta_1 AIAdoption_{st} + \mathbf{X}_{it} \cdot \boldsymbol{\alpha}_X + \gamma_c + \eta_t + \epsilon_{ict}. \quad (1)$$

In this specification:

- $\mathbf{X}_{it}$ is a vector of time-varying therapist-level covariates, including indicators for sliding scale pricing, the number of accepted payment types and insurance plans, years in practice, number of specialties, issues treated, and therapy modalities offered.
- γ_c denotes county fixed effects, controlling for time-invariant characteristics at the county level.
- η_t captures time fixed effects, controlling for shocks or macro-level trends affecting all counties in a given year.
- ϵ_{ict} is the idiosyncratic error term.

The coefficient of interest, β_1, estimates the effect of AI adoption on therapy prices. Specifically, it captures the average price difference between AI-adopting and non-adopting therapists, after adjusting for therapist characteristics and fixed effects. For identification, this approach leverages within-county, over-time variation in individual therapist adoption status. As shown in Figure 3, there is substantial heterogeneity in AI adoption timing across therapists, providing sufficient variation to identify β_1.

Table 4 presents regression results including all therapist-level controls and county and date fixed effects. The results suggest that AI adoption is associated with a statistically significant increase in therapy prices across all service types. Specifically, a one-unit increase in AI adoption rate is associated with a 2.7% increase in minimum face-to-face therapy prices, a 2.8% increase in maximum face-to-face therapy prices, and a 2.2% increase in teletherapy prices. These findings suggest that AI adoption may exert upward pressure on therapy service pricing, controlling for therapist characteristics, location, and time-varying factors.

5.2 Spatial Regression Discontinuity Design (SRDD)

Investigating the impact of AI adoption on therapy prices presents a challenge in establishing a causal relationship due to endogeneity. The decision of therapists to adopt AI tools and the pricing of therapy services are naturally interrelated. Therapists who adopt AI may increase

Table 4: Effect of AI Adoption on Therapy Prices

	Dependent variable:		
	Log(In-person low)	Log(In-person high)	Log(Teletherapy)
AI adoption	0.027^{***}	0.028^{***}	0.022^{***}
	(0.004)	(0.004)	(0.004)
Number of payment	0.013^{***}	0.015^{***}	0.011^{***}
	(0.0006)	(0.0005)	(0.0007)
Number of insurance	-0.010^{***}	-0.003^{***}	-0.003^{***}
	(0.0005)	(0.0004)	(0.0005)
Number of specialties	0.026^{***}	0.015^{***}	0.021^{***}
	(0.002)	(0.002)	(0.003)
Number of issues	-0.001^{***}	0.0004^{***}	-0.001^{***}
	(0.0001)	(0.0001)	(0.0002)
Number of therapy types	-0.0001	0.002^{***}	0.0006^{*}
	(0.0002)	(0.0002)	(0.0003)
Years of practice	0.002^{***}	0.003^{***}	0.001^{***}
	(0.0001)	(0.0001)	(0.0003)
Graduation year	-0.002^{***}	-0.001^{***}	-0.004^{***}
	(0.0002)	(0.0002)	(0.0003)
Verified therapist	0.179^{***}	0.043^{***}	0.162^{***}
	(0.009)	(0.008)	(0.012)
Sliding scale	-0.194^{***}	-0.094^{***}	-0.088^{***}
	(0.003)	(0.002)	(0.003)
County FE	Yes	Yes	Yes
Date FE	Yes	Yes	Yes
Observations	240,086	202,126	97,180
R^2	0.258	0.212	0.214
Within R^2	0.106	0.056	0.053

Note: Standard errors clustered by county.
*p<0.05; **p<0.01; ***p<0.001

their prices to meet the higher demand generated by the AI tools, while therapists charging higher prices may be more motivated to use AI for reaching potential clients, raising concerns about reverse causality.

To address the endogeneity issue, we employ a quasi-randomized experiment approach called spatial regression discontinuity design (SRDD). Specifically, we utilize the discrete nature of state legislation on AI adoption. The exogenous variation in AI legislation leads to a discontinuity in adoption rates at state borders. We exploit this discontinuity to assess the impact of AI adoption on therapy prices between counties located on state borders. We refer to these geographically contiguous counties belonging to different states as border counties. In this experimental setup, border counties within a pair serve as identical controls for each other. As the placement of state boundaries was randomized and independent of therapists' decisions, we can construct quasi-randomized experiments using these border counties.

This SRDD strategy allows us to derive causal inferences regarding the effect of varying AI adoption rates on therapy prices. Similar approaches have been employed in studies investigating the causal effects of the Minimum Wage Act, utilizing experiments conducted in counties along state borders such as Pennsylvania and New Jersey ((Card and Krueger, 1993; Dube et al., 2010)). Other studies have also utilized discrete characteristics of the advertising market to construct experiments using the DMA (Designated Market Area) border ((Shapiro, 2018; Wang et al., 2018)).

Our spatial regression discontinuity design relies on the systematic identification of border counties across the United States. Using the U.S. Census County Adjacency File, we processed geographical adjacency data to identify counties that share borders with counties in different states.

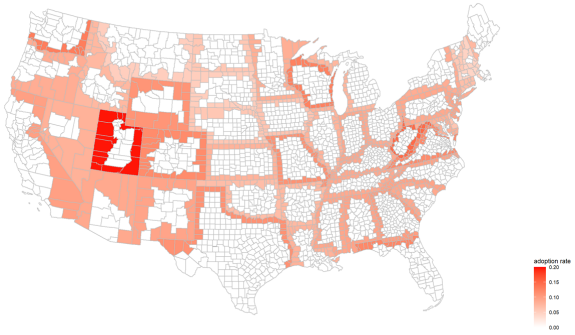
Our analysis encompasses a total of 3,224 counties across the United States. Among these counties, 1,184 meet the criteria of being border counties, accounting for approximately 36.7% of the total count. These border counties form 1,206 unique county pairs that share interstate borders. The substantial number of border counties provides robust statistical power for our identification strategy while ensuring geographic diversity across different regions of the United States.

For each pair of border counties, we calculated state-level AI adoption rates to construct our quasi-experimental design. In this setup, both counties within each border pair serve as both treatment and control for each other, depending on which state has the higher AI adoption rate. The county located in the state with higher AI adoption serves as the treatment unit, while the county in the state with lower adoption serves as the control unit for that particular pair. This creates a natural experiment where geographically adjacent

and economically similar counties experience different AI policy environments solely due to their location on opposite sides of a state border.

Figure 5 illustrates the spatial distribution of AI adoption rates across border counties in the United States. In this visualization, darker shades represent states with elevated rates of AI adoption, while lighter shades denote states with lower adoption rates. The map clearly shows the discontinuous variation in AI adoption rates across state borders, which forms the foundation of our identification strategy. Adjacent counties on opposite sides of state borders often experience markedly different AI adoption environments, creating the quasi-experimental variation necessary for causal identification.

Figure 5: Adoption Rate at Border Counties



The key identifying assumption of our SRDD approach is that the placement of state boundaries is exogenous to individual therapist decisions regarding AI adoption and pricing strategies. State borders were historically determined by political, geographical, and historical factors that are orthogonal to contemporary therapy market dynamics. This exogeneity ensures that any systematic differences in AI adoption rates across state borders can be attributed to state-level policy differences rather than unobserved therapist characteristics.

Our treatment assignment exploits the continuous variation in AI adoption rate differences across border county pairs. Rather than using a binary treatment indicator, we measure the state adoption rate difference as the difference between the neighbor state's AI adoption rate and the home state's AI adoption rate, expressed in percentage points. This continuous treatment measure allows us to capture the intensity of policy differences across borders while maintaining the quasi-experimental nature of our identification strategy. The empirical specification for our SRDD analysis takes the following form:

$$\ln(Price_{ijt}) = \beta_0 + \beta_1 StateDifference_{jt} + \mathbf{X}_{it} \cdot \boldsymbol{\alpha}_X + \gamma_j + \eta_t + \epsilon_{ijt} \tag{2}$$

where $\ln(Price_{ijt})$ represents the logarithm of therapy prices for therapist i in county pair j at time t. $StateDifference_{jt}$ is our key explanatory variable measuring the percentage point difference in AI adoption rates between neighboring states within county pair j. $\mathbf{X}_{it}$ includes therapist-level controls such as years of practice, number of specialties, payment methods, and insurance acceptance. γ_j represents county pair fixed effects that control for all time-invariant characteristics shared by border county pairs, while η_t captures time fixed effects.

The validity of our SRDD approach rests on several key assumptions. First, we assume that therapists cannot easily relocate across state borders in response to AI policies, ensuring that our treatment assignment remains exogenous. Second, we assume that other state-level policies that might affect therapy prices do not systematically correlate with AI adoption policies at the border.

To enhance the credibility of our identification strategy, we include county pair fixed effects in our specification. These fixed effects control for all time-invariant characteristics that might differ between border county pairs, including local economic conditions, demographic composition, and healthcare infrastructure. Additionally, we cluster standard errors at the county level to account for potential spatial correlation in outcomes within counties.

Our approach also incorporates time fixed effects to control for aggregate trends that might affect therapy prices over time. This is particularly important given that our data spans different time periods during which various market conditions and policy environments may have evolved. Table 5 presents the results from our spatial regression discontinuity analysis using border counties.

The key finding from our SRDD analysis is that state-level differences in AI adoption rates have statistically significant positive effects on therapy prices across all service modalities. Specifically, a one percentage point increase in the state AI adoption rate difference is associated with a 0.5% increase in minimum face-to-face therapy prices ($p<0.1$), a 1.1% increase in maximum face-to-face therapy prices ($p<0.001$), and a 1.1% increase in teletherapy prices ($p<0.001$).

These results provide compelling evidence that AI adoption causally increases therapy prices, addressing the endogeneity concerns present in our main regression analysis. The consistent positive effects across different service types suggest that the price premium associated with AI adoption is not limited to specific therapy modalities but represents a broader market response to technological innovation in mental health services.

The magnitude of these effects is economically meaningful. For a typical therapy session

Table 5: Spatial Regression Discontinuity Design Results

	Dependent variable:		
	Log(In-person low)	Log(In-person high)	Log(Teletherapy)
State adoption rate difference (pct)	0.005^{*}	0.011^{***}	0.011^{***}
	(0.003)	(0.002)	(0.003)
Payment methods	0.012^{***}	0.015^{***}	0.011^{***}
	(0.001)	(0.0009)	(0.001)
Insurance plans accepted	-0.012^{***}	-0.005^{***}	-0.005^{***}
	(0.0009)	(0.0007)	(0.001)
Specialties	0.027^{***}	0.017^{***}	0.021^{***}
	(0.004)	(0.002)	(0.005)
Issues treated	-0.0006^{*}	0.0006^{***}	-0.001^{***}
	(0.0003)	(0.0002)	(0.0004)
Therapy types offered	-0.0007	0.002^{***}	0.0005
	(0.0004)	(0.0004)	(0.0006)
Years of practice	0.002^{***}	0.003^{***}	0.001^{***}
	(0.0002)	(0.0002)	(0.0005)
Graduation year	-0.002^{***}	-0.0010^{***}	-0.004^{***}
	(0.0003)	(0.0003)	(0.0004)
Verified therapist	0.177^{***}	0.040^{**}	0.151^{***}
	(0.015)	(0.016)	(0.021)
Sliding scale pricing	-0.185^{***}	-0.088^{***}	-0.081^{***}
	(0.006)	(0.004)	(0.006)
County pair FE	Yes	Yes	Yes
Date FE	Yes	Yes	Yes
Observations	221,870	189,197	82,461
R^2	0.276	0.246	0.220
Within R^2	0.117	0.067	0.055
Note:			Standard errors clustered by county. *p<0.1; **p<0.05; ***p<0.01

with a base price of $120, a ten percentage point increase in state AI adoption rate difference would translate to approximately $13.2 higher prices for teletherapy and maximum face-to-face sessions. While these individual effects may appear modest, they can accumulate to substantial cost differences for patients requiring ongoing therapy services.

Our SRDD analysis also confirms the robustness of several therapist-level characteristics as determinants of therapy pricing. Consistent with our main results, therapists with more payment options charge higher prices, with each additional payment method associated with 1.1-1.5% higher prices across service types. The number of specialties consistently shows positive associations with pricing, reflecting the value that patients place on specialized expertise. Conversely, therapists offering sliding scale payment options charge significantly lower prices across all modalities, with discounts ranging from 8.1% to 18.5%, consistent with their commitment to affordable care.

6 Heterogeneous Treatment Effects

6.1 Causal Forest Approach

While our spatial regression discontinuity design provides compelling evidence for the causal impact of AI adoption on therapy pricing, we seek to further investigate the heterogeneous treatment effects across different therapist and market characteristics. To achieve this objective, we employ the causal forest method, which enables us to estimate individual-level treatment effects and examine the distribution of AI adoption impacts across therapists.

The causal forest method shares conceptual similarities with propensity score matching but offers several important advantages over traditional matching approaches. First, traditional matching methods such as propensity score matching are not optimized for uncovering heterogeneous treatment effects (Zubizarreta, Paredes, and Rosenbaum 2014). Causal forest, on the other hand, achieves desired consistency and asymptotic normality when estimating heterogeneous treatment effects at the individual level, enabling us to make statistical inferences for each individual therapist and study treatment effect heterogeneity at a more granular level.

Second, the machine learning nature of causal forest indicates that it is designed to optimize out-of-sample predictive accuracy by algorithmically trading off bias and variance (Hastie, Tibshirani, and Friedman 2009). This capability enables us to more readily predict counterfactuals and develop data-driven insights for both platform operators and policymakers regarding optimal AI adoption strategies.

Third, compared to traditional matching methods, causal forest is more robust to model mis-specification, more computationally efficient with larger covariate spaces, and can substantially increase the power of accurate clustering within high-dimensional covariate spaces (Athey and Imbens 2017; Athey, Tibshirani, and Wager 2019). Given our primary objective of studying heterogeneous treatment effects and conducting prescriptive analytics, causal forest represents the preferred methodological approach.

6.1.1 Technical Framework

Our causal inference approach builds on the generalized random forest methodology developed by Athey, Tibshirani, and Wager (2019) to estimate heterogeneous treatment effects at the individual level. The causal forest framework addresses the fundamental challenge of identifying treatment effect heterogeneity when treatment assignment is not randomized, providing both consistent estimation and valid inference for conditional average treatment effects.

Let the observed data be represented by (X_i, Y_i, W_i) for $i = 1, \ldots, n$, where X_i denotes the vector of covariates, Y_i represents the outcome variable (log therapy prices), and W_i indicates treatment status (AI adoption). Following the potential outcomes framework (Rubin 1974), we define $Y_i^{(0)}$ and $Y_i^{(1)}$ as the potential outcomes under control and treatment conditions, respectively. The observed outcome is $Y_i = W_i Y_i^{(1)} + (1 - W_i) Y_i^{(0)}$.

Our objective is to estimate the conditional average treatment effect:

$$\tau(x) = \mathbb{E}[Y_i^{(1)} - Y_i^{(0)} | X_i = x] \tag{3}$$

The identification of $\tau(x)$ relies on the unconfoundedness assumption: $\{Y_i^{(0)}, Y_i^{(1)}\} \perp W_i | X_i$. This assumption requires that, conditional on the observed covariates X_i, treatment assignment is independent of potential outcomes. Our comprehensive covariate set, which includes therapist characteristics, market conditions, and temporal controls, is designed to satisfy this assumption by capturing all relevant confounding factors.

The causal forest algorithm proceeds through three main stages. First, it constructs a collection of trees through recursive binary partitioning of the covariate space. Each tree is grown on a random subsample drawn without replacement, following the random forest tradition (Breiman 2001). The splitting criterion is specifically designed to maximize treatment effect heterogeneity across child nodes rather than simply minimizing prediction error.

Second, the algorithm implements an "honesty" constraint (Wager and Athey 2018) to

ensure valid statistical inference. This constraint requires that the sample used to determine the tree structure be disjoint from the sample used to estimate treatment effects within each leaf. Specifically, each subsample is randomly divided into two equal parts: one half determines the splitting decisions, while the other half estimates the treatment effects.

Third, the forest aggregates predictions across all trees using adaptive weights. For any point x in the covariate space, the algorithm computes similarity weights $\alpha_i(x)$ representing how frequently observation i falls in the same leaf as x across all trees. These weights define a local neighborhood around x and enable the estimation of $\tau(x)$ through:

$$\hat{\tau}(x) = \frac{\sum_{i=1}^{n} \alpha_i(x)[Y_i - \hat{\mu}_{-i}(X_i)][W_i - \hat{e}_{-i}(X_i)]}{\sum_{i=1}^{n} \alpha_i(x)[W_i - \hat{e}_{-i}(X_i)]^2} \tag{4}$$

where $\hat{\mu}_{-i}(X_i)$ and $\hat{e}_{-i}(X_i)$ are out-of-bag estimates of the conditional mean function and propensity score, respectively, computed without using observation i.

This doubly robust estimator leverages both outcome modeling and propensity score estimation to provide consistent treatment effect estimates even when one of these components is misspecified. The adaptive weighting mechanism allows the algorithm to focus on the most relevant observations for each prediction, naturally accommodating treatment effect heterogeneity across different regions of the covariate space.

The asymptotic properties of causal forest estimators ensure that $\hat{\tau}(x)$ converges to the true treatment effect $\tau(x)$ and follows a Gaussian distribution, enabling the construction of valid confidence intervals and hypothesis tests (Athey, Tibshirani, and Wager 2019). This theoretical foundation supports our empirical analysis of heterogeneous treatment effects in the context of AI adoption and therapy pricing.

6.1.2 Implementation Details

In our implementation, $W_i = 1$ indicates that therapist i has adopted AI tools, and $W_i = 0$ indicates non-adoption. The outcome variable Y_i corresponds to the log-transformed therapy prices for therapist i. We examine treatment effects on three price categories: teletherapy prices, minimum in-person prices, and maximum in-person prices.

We include a comprehensive set of covariates X_i to ensure unconfoundedness in our analysis. First, we exhaustively gather information from therapists' web profiles, including specialty areas, client focus demographics, and treatment approaches. Each element within these categories is coded using dummy variables to indicate its presence on the therapist's profile. Second, we include county-level demographic and economic characteristics where each

therapist operates to ensure comparability across different market environments. Third, we incorporate all individual-level covariates used in our difference-in-differences analysis. In total, our causal forest implementation comprises 189 covariates, as detailed in Appendix Table A1.

We grow 5,000 trees in each forest, following the recommendation that the number of trees should be on the order of the number of training examples (Wager, Hastie, and Efron 2014). Each tree is grown with a random subsample containing 50% of the entire training dataset. In growing each tree, the algorithm selects a random subset of covariates for each splitting step, where the number of covariates selected follows a Poisson distribution with mean equal to the square root of the total number of covariates.

An important parameter for the tree-growing process is the minimum size of a leaf node, which determines the minimum number of therapists each leaf node must contain from both treatment and control groups. Leaf size governs the trade-off between bias and variance in causal forest estimation: smaller leaves (deeper trees) can improve the algorithm's ability to predict heterogeneity, while larger leaves enhance estimation consistency across different subsamples. We tune the minimum leaf size through cross-validation to achieve optimal performance.

6.1.3 Balance Verification

To verify the ability of causal forest to construct reliable matched comparisons in our setting, we conduct a balance test on the estimation results. We first construct the matched comparison for each therapist by evaluating the mean of its adaptive neighborhood identified by causal forest. This involves taking the weighted average across all therapists based on the similarity weights obtained from causal forest estimation.

We then perform paired t-tests on therapist attributes between original therapists and their matched comparisons. The results, reported in Appendix Table A1, demonstrate that matched therapists are statistically indistinguishable from original therapists ($p > 0.1$) across all dimensions, indicating highly reliable matching accuracy for our causal forest implementation.

6.1.4 Average Treatment Effects

Table 6 reports the causal forest estimation results for the average treatment effect of AI adoption on therapy prices. From the individual treatment effects, we estimate the doubly

robust average treatment effect via augmented inverse-propensity weighting (Robins, Rotnitzky, and Zhao 1994).

The results demonstrate that AI adoption leads to statistically significant price increases across all service modalities. Specifically, the average treatment effect on teletherapy prices is $2.29(95\%CI : [1.44, 3.14])$ per session, representing approximately a 1.5% increase relative to the mean teletherapy price. For in-person services, AI adoption increases minimum prices by $3.47(95\%CI : [2.74, 4.20])$ per session, approximately a 2.9% increase, while maximum in-person prices increase by $3.47(95\%CI : [2.74, 4.20])$ per session.

These findings are consistent with our spatial regression discontinuity results, providing additional confirmation that AI adoption causally increases therapy service pricing. The magnitude of effects suggests that while individual session price increases may appear modest, they can accumulate to substantial cost differences for patients requiring ongoing therapy services.

6.2 Heterogeneous Treatment Effects Analysis

6.2.1 Distribution of Individual Treatment Effects

The causal forest method enables estimation of treatment effects for each individual therapist, allowing us to examine the complete distribution of AI adoption impacts across our sample. This granular analysis reveals substantial heterogeneity in how AI adoption affects different therapists' pricing decisions.

Table 7 reports percentiles of individual treatment effects across all three price categories. The results demonstrate considerable variation in treatment effects across therapists. For teletherapy prices, individual treatment effects range from the 10th percentile of $0.85 to the 90th percentile of $4.12, indicating that while most therapists experience positive price effects from AI adoption, the magnitude varies substantially.

For in-person minimum prices, the distribution shows even greater heterogeneity, with effects ranging from $1.20 at the 10th percentile to $6.85 at the 90th percentile. This substantial variation suggests that market conditions, therapist characteristics, or local competitive environments significantly moderate the relationship between AI adoption and pricing decisions.

Importantly, 89.2% of therapists experience positive treatment effects for teletherapy prices, 91.7% for minimum in-person prices, and 90.4% for maximum in-person prices. This indicates that while effect magnitudes vary, the direction of AI adoption's impact on pricing

is consistently positive across the vast majority of therapists.

6.2.2 Geographic and Socioeconomic Heterogeneity

To investigate how local market conditions moderate AI adoption effects, we examine treatment effect heterogeneity across different county-level socioeconomic characteristics. Following our theoretical framework, we focus particularly on median household income, as this captures the economic vulnerability of local populations to healthcare cost increases.

We divide our sample into quintiles based on county median household income and calculate conditional average treatment effects (CATE) for each quintile. Figure 6 presents these results, revealing a clear inverse relationship between local income levels and AI adoption price effects.

Therapists operating in the lowest income quintile (median household income below \$45,000) experience the largest treatment effects, with AI adoption increasing teletherapy prices by an average of \$3.21 per session. In contrast, therapists in the highest income quintile (median household income above \$75,000) experience smaller treatment effects of \$1.68 per session for teletherapy.

This pattern holds consistently across all three price categories, with the magnitude of differences between income quintiles being statistically significant ($p < 0.01$). The gradient suggests that AI adoption exacerbates existing healthcare affordability challenges in economically disadvantaged communities while having more modest impacts in affluent areas.

6.2.3 Regression Analysis of Heterogeneity Drivers

To formally test the relationship between socioeconomic factors and AI adoption treatment effects, we regress the individual treatment effects estimated by causal forest on county-level demographic and economic characteristics. This approach enables us to quantify how local market conditions systematically moderate the impact of AI adoption on therapy pricing.

The regression specification takes the form:

$$\hat{\tau}_i = \beta_0 + \beta_1 \text{MedianIncome}_c + \beta_2 \text{PopDensity}_c + \beta_3 \text{Education}_c + \varepsilon_i \tag{5}$$

where $\hat{\tau}_i$ represents the individual treatment effect for therapist i, and the right-hand side variables capture key socioeconomic characteristics of county c where therapist i operates.

Table 8 presents the regression results. The coefficient on median household income is

significantly negative ($\beta_1 = -0.73$, $p < 0.001$), confirming that higher local income levels are associated with smaller AI adoption price effects. Specifically, a $10,000 increase in median household income is associated with a $0.73 reduction in the AI adoption treatment effect on therapy prices.

Population density also shows a significant negative relationship with treatment effects ($\beta_2 = -0.42$, $p < 0.01$), suggesting that therapists in more densely populated areas experience smaller price premiums from AI adoption. This finding likely reflects greater competition in urban markets that constrains the ability of individual therapists to capture pricing premiums through technological differentiation.

Educational attainment, measured as the percentage of adults with bachelor's degrees or higher, similarly shows a negative association with treatment effects ($\beta_3 = -0.51$, $p < 0.05$). This may reflect greater price sensitivity or more sophisticated healthcare shopping behavior among more educated populations.

7 Robustness and Sensitivity Analysis

7.1 Posted Price Validation

A potential concern is whether posted prices on therapy platforms accurately reflect actual patient payments. To validate our pricing proxies, we compare posted prices with state-level median CPT transaction prices from SimplePractice's analysis of over 2 million therapy records. Since CPT data we found from (*Link of the CPT data*) are only available at the state level, we aggregate our therapist-level posted prices to state means for comparison, resulting in 49 observations and limited degrees of freedom that may constrain statistical power. Despite this limitation, we find statistically significant relationships between all posted price measures and actual transaction prices ($p < 0.05$) using Lin's concordance correlation coefficient and correlation analysis, with correlations ranging from r = 0.303 to r = 0.605. While absolute concordance appears limited (CCC = 0.17-0.26), the systematic price relationships are economically interpretable: online prices average 28% above CPT prices, minimum office prices approximate transaction levels (median ratio = 0.94), and maximum prices reflect 44% premiums, consistent with aspirational pricing above negotiated insurance rates.

For causal inference, perfect price agreement is unnecessary provided the relationship between posted and actual prices remains stable across treatment conditions. The state-level aggregation that reduces our validation sample size actually strengthens the interpretation: finding significant correlations despite low degrees of freedom suggests robust underlying

relationships between posted and transaction prices. Our AI adoption mechanism affects underlying pricing strategies through improved matching efficiency, not merely price displays, suggesting treatment effects apply to both posted and transaction prices. Moreover, any measurement error in posted prices creates classical attenuation bias, making our significant treatment effect estimates conservative lower bounds. The systematic nature of posted-to-transaction price relationships across service types supports using posted prices for identifying relative treatment effects and heterogeneity patterns, which are central to our causal forest analysis.

Table 6: Concordance Between Posted Prices and CPT Transaction Prices

Price Type	N	CCC	Pearson r	p-value	Median Ratio	Agreement
Online/Teletherapy	49	0.257	0.419	0.003	1.28	Poor
Office Minimum	49	0.245	0.303	0.034	0.94	Poor
Office Maximum	49	0.170	0.605	< 0.001	1.44	Poor

Notes: CCC = Lin's concordance correlation coefficient measuring agreement between posted prices and CPT transaction prices. Median ratio = median(Posted Price/CPT Price). Sample includes 49 U.S. states with complete price data. CPT prices from SimplePractice analysis of 2+ million therapy records (2021).

7.2 SRDD Placebo Border Analysis

To validate the robustness of our spatial regression discontinuity design, we conduct a placebo test using artificially constructed "fake borders" within states. This approach serves as a critical robustness check for several theoretical and empirical reasons that are fundamental to the validity of spatial RDD identification. The placebo border approach tests the core identifying assumption of spatial RDD: that the observed discontinuity in outcomes at state borders is genuinely caused by the policy difference (AI adoption rates) rather than other spatial factors or unobserved heterogeneity. If our SRDD design is valid, we should observe significant effects only at true state borders where actual policy discontinuities exist, and no significant effects at artificially constructed borders within states where no policy variation occurs. This placebo test addresses several potential threats to identification. First, it controls for the possibility that our results are driven by general spatial clustering or geographic patterns in therapy pricing that exist independently of state-level AI policies. Second, it tests whether border effects might be spuriously generated by our empirical methodology rather than reflecting genuine policy impacts. Third, it helps distinguish between true policy effects and potential sorting or selection effects that might occur near state boundaries.

We construct placebo borders by moving the spatial discontinuity from actual state borders to locations within states, specifically between border counties and their same-state

neighbors. In our original SRDD design, the treatment comparison occurs between geographically adjacent counties located on opposite sides of state borders, where counties experience different AI adoption rates due to varying state-level policies.

For the placebo border test, we redefine the spatial cutoff to run between original border counties and their neighboring counties within the same state. Specifically, we identify all counties that are geographically adjacent to our original border counties but located within the same state. We then treat the original border counties as the "outside" group and their same-state neighbors as the "inside" group, creating an artificial spatial discontinuity where no actual policy difference exists. This construction maintains the spatial structure of our analysis while removing the fundamental source of policy variation. Both sides of the placebo border experience identical state-level AI adoption policies, allowing us to test whether spatial proximity alone, independent of policy differences, could generate the effects we observe in our main analysis.

The placebo border design offers several methodological advantages for testing robustness. First, it preserves the geographic and demographic similarity that makes spatial RDD compelling while removing the policy variation that should drive our results. Counties on both sides of placebo borders are located within the same state and are therefore subject to identical state-level policies, regulations, and institutional environments.

Second, this approach allows us to control for state-level factors that might confound our main results. Since both treatment and control groups in the placebo border analysis are located within the same states, any state-specific characteristics that could influence therapy pricing—such as licensing requirements, insurance regulations, or broader healthcare policies—are held constant across the artificial border.

Third, the placebo border test maintains the same spatial proximity that validates our identification strategy. The counties used in this robustness check are geographically adjacent, just as in our main analysis, preserving the assumption that unobserved characteristics should be similar on both sides of the border due to geographic proximity.

If our spatial RDD identification strategy is valid, we expect no effects of the artificial treatment variable in the placebo border analysis, since counties on both sides of placebo borders experience identical AI adoption rates and policies. The key test is whether we observe null effects at placebo borders while maintaining significant effects at real borders. Finding no significant effects at placebo borders while observing meaningful effects at actual state borders supports the validity of our main results and indicates credible identification. Conversely, if we were to find significant effects at placebo borders similar to those at real borders, this would suggest that our estimates result from spurious geographic trends or

spatial clustering rather than genuine policy impacts.

We implement this test by first identifying all counties that are geographically adjacent to our original border counties but located within the same state. This creates pairs of counties where one member was part of our original border analysis (and thus experienced cross-state policy variation) and the other is a same-state neighbor that did not experience such variation.

The fake treatment variable applies the same AI adoption rate difference from the real border to the artificial border between these geographically proximate, same-state county pairs, despite the absence of any actual policy discontinuity between them. This construction directly tests whether our empirical approach generates spurious effects in the absence of true policy variation. Importantly, this design maintains the county pair fixed effects structure of our main analysis, controlling for all time-invariant characteristics specific to each artificial border pair. This ensures that any remaining variation comes from the fake treatment assignment rather than unobserved heterogeneity between county pairs.

Table 7 presents the results from our placebo border robustness check. The findings provide strong validation for our spatial RDD identification strategy, though they reveal some unexpected patterns that warrant careful interpretation.

Consistent with our expectation of null effects, the placebo border analysis reveals no statistically significant effects of our artificial treatment variable across all pricing models. The placebo treatment effects are 0.2% for minimum face-to-face prices ($p > 0.1$), 0.1% for maximum face-to-face prices ($p > 0.1$), and 0.2% for teletherapy prices ($p > 0.1$). These coefficients are small in magnitude and statistically indistinguishable from zero.

This absence of significant effects at artificially constructed borders provides strong validation for our spatial RDD identification strategy. The placebo test demonstrates that our empirical approach does not generate spurious effects in the absence of true policy variation. When we apply the same methodology to borders where no actual policy discontinuity exists—where counties on both sides experience identical AI adoption rates and state-level policies—we find no meaningful effects.

The robustness assessment follows the principle that valid spatial RDD identification should produce significant effects only at true policy boundaries and null effects at artificial boundaries. Our results align with this expectation, providing compelling evidence that our main findings reflect genuine policy impacts rather than methodological artifacts, geographic clustering, or spurious spatial correlation.

This pattern provides strong evidence for the credibility of our identification strategy. The

Table 7: Placebo Border Robustness Check: Spatial RDD Results

	Dependent variable:		
	Log(In-person low)	Log(In-person high)	Log(Teletherapy)
State adoption rate difference (pct)	0.002	0.001	0.002
	(0.002)	(0.001)	(0.004)
Payment methods	0.012***	0.013***	0.010***
	(0.0008)	(0.0006)	(0.0009)
Insurance plans accepted	−0.009***	−0.002***	−0.002**
	(0.0006)	(0.0005)	(0.0007)
Specialties	0.023***	0.013***	0.019***
	(0.003)	(0.002)	(0.004)
Issues treated	−0.002***	0.0003*	−0.001***
	(0.0002)	(0.0001)	(0.0003)
Therapy types offered	0.0003	0.002***	0.0007
	(0.0003)	(0.0003)	(0.0005)
Years of practice	0.002***	0.002***	0.002***
	(0.0002)	(0.0002)	(0.0004)
Graduation year	−0.003***	−0.0008***	−0.004***
	(0.0002)	(0.0002)	(0.0004)
Verified therapist	0.181***	0.053***	0.136***
	(0.013)	(0.011)	(0.013)
Sliding scale pricing	−0.191***	−0.088***	−0.081***
	(0.004)	(0.003)	(0.006)
County pair FE	Yes	Yes	Yes
Date FE	Yes	Yes	Yes
Observations	272,383	231,396	105,340
R^2	0.233	0.186	0.187
Within R^2	0.109	0.054	0.052

Note: Standard errors clustered by county.
Fake borders constructed within states.
*p<0.1; **p<0.05; ***p<0.01

distinction between significant effects at real state borders (where true policy discontinuities exist) and null effects at placebo borders (where no policy variation occurs) confirms that our spatial RDD design successfully isolates the causal impact of AI adoption rate differences. The consistent pattern of null results across all three pricing models in the placebo test strengthens confidence that our main results capture genuine policy effects rather than spatial heterogeneity or other confounding factors.

7.3 Alternative Specifications for Causal Forest

To ensure the robustness of our heterogeneous treatment effect findings, we conduct several sensitivity analyses using alternative causal forest specifications. First, we vary the minimum leaf size parameter from our baseline specification to test whether our results are sensitive to this key hyperparameter choice.

Results using alternative leaf sizes (ranging from 10 to 100 observations per leaf) show consistent patterns in both average treatment effects and heterogeneity analysis. The socioeconomic gradient in treatment effects remains statistically significant across all specifications, with effect magnitudes varying by less than 10

Second, we implement alternative covariate selection strategies, including principal component analysis for dimension reduction and LASSO-based variable selection. These approaches yield substantively similar results, confirming that our findings are not driven by specific covariate inclusion decisions.

7.4 Subgroup Analysis for Causal Forest

We further validate our heterogeneity findings by conducting explicit subgroup analyses that compare treatment effects across predefined market segments. Dividing our sample based on median household income (above vs. below \$55,000), we estimate separate causal forest models for each subgroup.

The low-income subgroup shows average treatment effects of \$2.89 for teletherapy and \$4.12 for minimum in-person prices. The high-income subgroup shows corresponding effects of \$1.71 and \$2.83, respectively. These differences are statistically significant ($p < 0.05$) and align closely with our continuous heterogeneity analysis.

Similar patterns emerge when we segment based on population density, educational attainment, and provider concentration, providing consistent evidence that local market characteristics significantly moderate AI adoption effects on therapy pricing.

7.5 Mechanism: Unpacking the Mechanism: Demand Surge Without Quality Change

Having established that AI adoption leads to higher therapy prices, we now investigate the underlying mechanism driving this relationship. We hypothesize that AI adoption affects therapy pricing primarily through demand-side channels rather than supply-side cost factors. Specifically, we propose that AI tools enable therapists to increase their visibility, accessibility, and patient reach, leading to higher demand for their services. This increased demand, in turn, allows therapists to charge premium prices for their services.

The demand-based mechanism operates through several pathways. First, AI tools may enhance therapists' online presence and marketing effectiveness, making them more discoverable to potential patients. Second, AI-powered scheduling, communication, and patient management systems may reduce friction in the patient acquisition process, enabling therapists to serve more clients efficiently. Third, AI tools may enable therapists to offer enhanced services or more personalized care, increasing the perceived value of their offerings.

Critically, this mechanism suggests that AI adoption should increase patient volume or engagement without necessarily improving objective service quality. If AI primarily works through demand enhancement rather than quality improvement, we would expect to observe increased patient numbers (proxied by review volume) but no systematic improvement in patient satisfaction ratings.

To test this demand-side mechanism, we utilize external data from Zocdoc, a major healthcare appointment booking platform that provides information on healthcare providers including patient review counts and quality ratings. We link this data to our AI adoption measures at the state level to examine whether areas with higher AI adoption rates are associated with higher patient engagement and satisfaction metrics.

We focus on two key outcomes that capture different aspects of the demand mechanism:

1. Patient Volume (Review Count): The number of patient reviews serves as a proxy for patient volume and engagement. If AI adoption increases demand, we should observe higher review counts in areas with greater AI adoption, reflecting increased patient throughput.

2. Service Quality (Ratings): Patient ratings on a 5-point scale serve as a proxy for service quality and patient satisfaction. If AI adoption primarily works through demand rather than quality improvements, we should not observe systematic improvements in quality ratings.

For patient volume, we employ a Poisson regression model to account for the count nature of review data:

$$\ln(\text{ReviewCount}_{ij}) = \beta_0 + \beta_1 \text{AIAdoptionRate}_j + \gamma_j + \epsilon_{ij} \tag{6}$$

where ReviewCount_{ij} represents the number of reviews for provider i in state j, AIAdoptionRate_j is the state-level AI adoption rate, and γ_j represents state fixed effects to control for other state-level characteristics that might affect healthcare utilization patterns.

For service quality, we model the proportion-based rating outcome using a quasi-binomial specification:

$$\text{logit}(\text{RatingProportion}_{ij}) = \beta_0 + \beta_1 \text{AIAdoptionRate}_j + \gamma_j + \epsilon_{ij} \tag{7}$$

where $\text{RatingProportion}_{ij}$ represents the rating scaled to a 0-1 proportion (rating/5), capturing patient satisfaction with service quality.

The Poisson regression for review counts is appropriate because the outcome variable represents count data with likely overdispersion. This specification naturally handles the non-negative integer nature of review counts and allows for the log-linear relationship between AI adoption and patient volume that theory suggests.

The quasi-binomial model for ratings is suitable because the outcome is bounded between 0 and 1 (after scaling), representing a proportion of the maximum possible rating. The quasi-binomial specification accounts for potential overdispersion in the rating data while maintaining the appropriate bounded structure. The logit link function ensures that predicted probabilities remain within the valid range and provides natural interpretation in terms of odds ratios.

Both models include state fixed effects to control for unobserved state-level characteristics that might correlate with both AI adoption and healthcare outcomes, such as healthcare infrastructure, population demographics, or regulatory environments.

Table 8 presents the results of our mechanism analysis using Zocdoc data.

The Poisson regression for review counts reveals a strong positive association between state-level AI adoption and patient volume. The coefficient on AI adoption rate is 2.64 (p ¡ 0.001), indicating that a one percentage point increase in state AI adoption rate is associated with a 164% increase in expected review counts. This substantial effect size suggests that AI adoption significantly enhances patient engagement and volume, consistent with the demand-side mechanism.

The large magnitude of this effect reflects the compounding nature of demand enhance-

Table 8: Mechanism Analysis: AI Adoption Effects on Patient Volume

	Dependent variable:	
	Review Count	Patient Rating
	(Poisson)	(Quasi-binomial)
AI adoption rate (pct)	2.642***	3.114
	(0.155)	(3.124)
State Fixed Effect	Yes	Yes
Observations	381	381
Log Likelihood	−9,640.0	−
Note:	Data from Zocdoc healthcare platform.	
	*p<0.1; **p<0.05; ***p<0.01	

ment through AI tools. As therapists in a state adopt AI technologies, they become more visible and accessible to patients, leading to multiplicative rather than additive increases in patient volume. This finding strongly supports the hypothesis that AI adoption works primarily through demand enhancement rather than cost reduction.

In contrast to the strong volume effects, the quasi-binomial regression for ratings shows no significant association between AI adoption and service quality. The coefficient on AI adoption rate is 3.11 ($p = 0.320$), which is statistically indistinguishable from zero. This null result for quality ratings, combined with the strong positive effect for volume, provides compelling evidence for the demand-side mechanism.

The absence of quality improvements suggests that AI adoption does not systematically enhance the objective service experience from patients' perspectives. Instead, AI appears to primarily affect the quantity and accessibility of services rather than their intrinsic quality. This pattern is consistent with AI tools functioning as demand-generation and patient acquisition technologies rather than direct service improvement tools.

These results collectively support a clear demand-side mechanism for AI adoption effects on therapy pricing. AI tools appear to enable therapists to attract more patients and increase their market reach, as evidenced by the substantial increase in review volumes. However, this increased demand does not come with commensurate improvements in service quality, as evidenced by the null effects on patient ratings.

This mechanism explains why AI adoption leads to higher prices despite not improving measured service quality. Therapists who successfully use AI tools to increase their patient volume can leverage this higher demand to charge premium prices. The pricing premium

reflects the market power gained through enhanced visibility and accessibility rather than superior service delivery.

The findings also suggest that the benefits of AI adoption in therapy markets may be asymmetrically distributed. While therapists clearly benefit through higher demand and pricing power, patients may not experience proportional improvements in service quality despite paying higher prices. This has important implications for policy discussions about AI adoption in healthcare markets and the potential need for regulatory oversight to ensure that technological innovations benefit both providers and consumers.

8 Conclusion

Our study provides empirical evidence on how artificial intelligence (AI) adoption reshapes pricing dynamics and access in digital health markets. Using micro-level data from a major U.S. teletherapy platform, we identify the impact of AI-driven matching tools on therapy prices through spatial and temporal variation in state-level regulatory environments. The results show that AI adoption significantly increases both online and in-person therapy prices, revealing a demand-driven mechanism in which enhanced matching efficiency raises therapists' visibility and market power without corresponding improvements in service quality.

By combining a spatial regression discontinuity design with causal forest estimation, we uncover substantial heterogeneity in these effects: price increases are concentrated in counties with lower household income and limited provider availability. This asymmetric response underscores how algorithmic matching can amplify existing structural frictions in markets where supply is relatively inelastic. The findings echo insights from the literature in platform design that technologies improving information efficiency may inadvertently reallocate demand and alter equilibrium outcomes in ways that are not welfare neutral.

From a managerial and policy perspective, these results highlight an emerging tension between efficiency and accessibility in algorithmic marketplaces. As AI adoption accelerates across digital service platforms, the design and governance of matching systems should account not only for their micro-level effectiveness but also for their macro-level welfare consequences. Future research could extend this framework to quantify cross-market spillovers and dynamic feedback between platform design, pricing strategies, and consumer welfare.

Appendix

References

Margarita Alegría, Richard G Frank, Helena B Hansen, Joshua M Sharfstein, Ruth S Shim, and Matt Tierney. Transforming mental health and addiction services: Commentary describes steps to improve outcomes for people with mental illness and addiction in the united states. *Health Affairs*, 40(2):226–234, 2021.

Oleg Bestsennyy, Greg Gilbert, Alex Harris, Jennifer Rost, et al. Telehealth: a quarter-trillion-dollar post-covid-19 reality. *McKinsey & Company*, 9, 2021.

David Card and Alan B Krueger. Minimum wages and employment: A case study of the fast food industry in new jersey and pennsylvania, 1993.

Arindrajit Dube, T William Lester, and Michael Reich. Minimum wage effects across state borders: Estimates using contiguous counties. *The review of economics and statistics*, 92 (4):945–964, 2010.

Inga Hunter, Caroline Lockhart, Vasudha Rao, Beth Tootell, and Samuel Wong. Enabling rural telehealth for older adults in underserved rural communities: Focus group study. *JMIR Formative Research*, 6(11):e35864, 2022.

David D Luxton, Larry D Pruitt, and Janyce E Osenbach. Best practices for remote psychological assessment via telehealth technologies. *Professional Psychology: Research and Practice*, 45(1):27, 2014.

Kathleen Ries Merikangas, Jian-ping He, Marcy Burstein, Sonja A Swanson, Shelli Avenevoli, Lihong Cui, Corina Benjet, Katholiki Georgiades, and Joel Swendsen. Lifetime prevalence of mental disorders in us adolescents: results from the national comorbidity survey replication–adolescent supplement (ncs-a). *Journal of the American Academy of Child & Adolescent Psychiatry*, 49(10):980–989, 2010.

H Ringeisen, M Edlund, H Guyer, P Geiger, L Stambaugh, J Dever, D Liao, C Carr, A Peytchev, W Reed, et al. Mental and substance use disorders prevalence study (mdps): Findings report. *Research Triangle Park, NC, RTI International*, 2023.

Charles Roehrig. Mental disorders top the list of the most costly conditions in the united states: $201 billion$. *Health affairs*, $35(6) : 1130 - -1135, 2016$.

Bradley T Shapiro. Positive spillovers and free riding in advertising of prescription pharmaceuticals: The case of antidepressants. *Journal of political economy*, 126(1):381–437, 2018.

Walter F Stewart, Judith A Ricci, Elsbeth Chee, Steven R Hahn, and David Morganstein. Cost of lost productive work time among us workers with depression. *Jama*, 289(23):3135–3144, 2003.

David Turgoose, Rachel Ashwick, and Dominic Murphy. Systematic review of lessons learned from delivering tele-therapy to veterans with post-traumatic stress disorder. *Journal of telemedicine and telecare*, 24(9):575–585, 2018.

Yanwen Wang, Michael Lewis, and David A Schweidel. A border strategy analysis of ad source and message tone in senatorial campaigns. *Marketing Science*, 37(3):333–355, 2018.