Can AI Distort Human Capital?*

Xuelin Li† **Meizi Zhou**‡

This version: January, 1st 2024

Abstract

We document that interactions with AI can distort the evolution of human capital in the context of opioid prescriptions. Physicians in our sample adopted electronic health record software from a list of certified companies in 2011. Between 2016 and Spring 2019, one company secretly embedded a biased AI function in its clinical decision support system to promote the sales of extended-release opioids. Affected physicians not only increased opioid claims relative to the control group during the treatment window but also maintained a higher propensity for prescriptions even after the removal of the biased function. This long-term distortion of human capital relies on the unconsciousness of AI biases and does not occur following other explicit promotions, such as pharmaceutical detailing payments. Experience with opioids, along with caution regarding patient risks, mitigates the distortion.

Keywords: artificial intelligence, human capital, clinical decision support system, opioid

*We are grateful to Carri Chan, Hong Luo, Laura Veldkamp for helpful comments and suggestions.

†Columbia University, Graduate School of Business - xuelin.li@columbia.edu

‡Boston University, Questrom School of Business - mzzhou@bu.edu

1 Introduction

The advancement of artificial intelligence (AI) has sparked a heated debate on its implications for the labor market. Automation of tasks previously performed by humans implies a reduced demand for labor in production, creating the job displacement effect. However, new technology also increases productivity and reinstates labor into newly developed tasks (Acemoglu and Restrepo, 2018, 2020). Instead, this paper focuses on the development of human capital through interactions with AI. With recent achievements in machine learning techniques, the human decision-making process is nowadays commonly assisted by AI predictions (Agrawal et al., 2019). We argue that human capital will adjust through collaboration with AI: these predictions will not only affect the current decision, but also alter workers' habits when similar tasks emerge in the future.

Empirically documenting this effect presents several challenges. First, pinning down the specific influences of AI on human behaviors is difficult due to substantial heterogeneity in human-AI interactions across tasks, labors, and algorithms (e.g Kleinberg et al., 2018; Angelova et al., 2023). Even in similar tasks, a single worker may face various recommended decisions across scenarios, making it challenging to identify alterations in habits consistently. Second, an essential step in this study is to obtain the long-term counterfactual of worker behaviors without further AI influence after their initial interaction. However, workers endogenously adopt and abandon AIs based on their recent experiences. As AI becomes a "general-purpose technology" (Bresnahan and Trajtenberg, 1995), AI adoption is almost an irreversible decision and workers rarely suspend usage, particularly due to exogenous reasons.

We address these challenges in the context of opioid prescriptions by physicians assisted with clinical decision support systems (CDSS). The background is that the Affordable Care Act created financial incentives for physicians to use electronic health record (EHR) software from a list of certified companies in 2011. However, from July 2016 to the Spring of 2019, one major company, Practice Fusion, secretly embedded a CDSS pain alert to promote the sales of Purdue Pharma's extended-release opioids (EROs) in an "unbranded effort." This setting provides the following empirical advantages. First, the AI assistance was manipulated to unanimously recommend opioid prescriptions regardless of medical necessity and poten-

tial abuse risks. This homogeneity allows us to measure an intuitive and single-direction impact on human behaviors. Second, our identification strategy follows a simple difference-in-differences (DID) design, leveraging physician-level annual opioid prescription data from 2013 to 2021. Treated physicians adopted Practice Fusion in 2011, and we can compare their behaviors to other incentive program participants after the shock. Most importantly, our sample spans the period (2019 – 2021) when the bias alert was removed so that we can evaluate the long-term distortions on human capital.

We first confirm that in the treatment window from 2016 to 2018, physicians having adopted Practice Fusion exhibit a significantly higher propensity of prescribing long-acting (LA) opioids, i.e., the category encompassing all Purdue Pharma's EROs, compared to the control group. We characterize this propensity through various measures, including number of claims, costs of drugs, days of supply and the fraction of LA opioids over all types of opioid claims. The coefficients are not only highly statistically significant (t-statistics from 3.6 to 6.4) but also economically substantial. Affected physicians increase their annual number of LA claims by 7.6%, generating an aggregated cost of approximately $3.14 million every year in our sample. These results imply that even professionals with expertise rely on AI recommendations and cannot distinguish intentional manipulations embedded in the system.

If the previous findings were purely driven by the transient manipulation, we would not observe any significant differences between the treatment and control groups after the removal of the AI manipulation. However, in the post-treatment window from 2019 to 2021, we document evidence of human capital distortion across all different measures: affected physicians continue to prescribe 12.4% more LA opioid claims relative to the control group. We then estimate the coefficient dynamics across all calendar years to characterize the time series of impacts. Supporting the parallel trend assumption, the two groups exhibit no significant differences before 2015. When the pain alert was active, affected physicians gradually increased LA opioid usage over year. After 2019, this trend is flattened but the treatment group still maintain a significantly higher propensity for opioid prescriptions. This distortion not only costs Medicare Part D $5.2 million every year in our sample, but also generate long-term real healthcare consequences for beneficiaries. For each 3-digit zip code area, a 1% increase in the number of Practice fusion adopters is associated with 0.48% more

Medicare beneficiaries requiring treatment for opioid abuse.

AI manipulation has the potential to distort human capital beyond the scale of traditional financial incentives. We argue that the unawareness of AI manipulation plays a key role in habit formation, given that Purdue Pharma's public image became increasingly negative in this period. To show this, we focus on the alternative group of physicians having financial relationships with Purdue Pharma, who explicitly received in-kind payments from the company. Purdue Pharma suspended these detailing activities also in 2019, coinciding with the removal of the pain alert. In stark contrast, recipients exhibit a significantly lower volume of LA claims compared to the control group post-payment suspension, displaying no long-lasting persistence in prescriptions. This finding aligns with Purdue Pharma's motivation for contracting with Practice Fusion, claiming that direct EROs promotions were hindered by providers' political pressure. Instead, physicians are unconscious of of biases in AI predictions and perceive the recommendations as optimally generated from real data. Through this interaction, they gradually update more favorable beliefs about the benefit-risk profile of EROs over time, leading to an inclination towards long-term prescriptions.

Lastly, we hypothesize two channels through which affected physicians are more robust to potential AI manipulations. First, physicians with more experience, particularly in LA opioid usage, are more likely to make decisions based on their expertise. Second, physicians who interact with patients of higher risks are more cautious about the potential side-effects. We validate these hypotheses through heterogeneity tests based on the ex-ante physician features before the shock. In both cases, the distortion driven by the CDSS is attenuated.

This paper closely relates to the literature on the labor market implications of artificial intelligence, including but not limited to Acemoglu and Restrepo (2018), Autor and Salomons (2018), Acemoglu and Restrepo (2020), and Acemoglu and Restrepo (2022). Instead of focusing on labor share or wage, we study the development of human capital through AI interactions. On this end, existing literature focuses on worker skills across various tasks that are exposed to AI replacements or compatible with AI capacities (Brynjolfsson et al., 2018; Felten et al., 2018; Webb, 2019; Acemoglu et al., 2022). Our paper differs as a "within-task" analysis that studies how working with AI will change the long-term performance of professionals with specific skills. Our empirical results echo the theoretical prediction in

Acemoglu (2021). He argues that AI assistance will diminish workers' understanding of automated tasks, which in turn reduces the productivity in their specialized fields requiring human judgements.

Our work also adds to the literature of understanding the role of information technologies in the healthcare industry. For discussion and overview, see Bronsoler et al. (2022) and Dranove and Garthwaite (2022). Examples of CDSS benefits include reducing adverse medical events (Hydari et al., 2019), attenuating racial biases in healthcare delivery (Ganju et al., 2020), and avoiding high-cost orders (Doyle et al., 2019). Closely related is Agarwal et al. (2023). The authors document that professional radiologists do not fully capitalize on the potential gains from AI assistance due to deviations in Bayesian belief updating. Radiologists may underweight the AI's information relative to their own. In their work, signals acquired by professionals are invariant to the provision of AI assistance. Our work complements theirs by documenting how experience with AI can change the long-term beliefs of professionals, which can dynamically affect the deviations documented in their paper.

Lastly, this paper falls into the broad literature comparing human decisions to algorithm-based recommendations in tasks such as diagnosing (Mullainathan and Obermeyer, 2022), judging (Kleinberg et al., 2018; Angelova et al., 2023), asset management (D'Acunto et al., 2019; Rossi and Utkus, 2020), and financial analysts (Cao et al., 2021; Coleman et al., 2022; Grennan and Michaely, 2020). In these settings, algorithms may unintentionally amplify biases embedded in unrepresentative training samples, creating fairness concerns (Cowgill and Tucker, 2019; Fuster et al., 2022). We document the costs of AI from a different perspective: firms can deliberately manipulate the predictions, exploiting humans' trust on AI, to achieve specific goals. Even professionals with expertise will have a hard time figuring out the manipulation.

2 Institutional Details and Data

2.1 Practice Fusion's CDSS

Practice Fusion was an EHR company founded in 2005. At one point, it was named the top EHR for customer satisfaction among primary care providers and ranked No. 1 for value among ambulatory professionals. Investors including Kleiner Perkins Caufield & Byers made a $70 million Series D investment in September 2013. This new round valued the company at around $700 million, making it one of the largest digital health startups at the time.[1]

Beginning in Fall 2013, Practice Fusion solicited remuneration and negotiated with a pharmaceutical company "Pharma Co.X," later identified as Purdue Pharma, in exchange for creating and embedding a pain alert in its CDSS.[2] The purpose is to increase prescriptions for Purdue Pharma's three ERO products. During their communication, Purdue Pharma expressed concern that "providers are hesitant about using high dosages to combat pain for a variety of reasons, mostly political pressure." In response, Practice Fusion proposed developing a pain alert to initiate ERO products and embedding Purdue Pharma's ERO products as treatment options in an "unbranded effort."

The two companies signed the official contracts on March 1, 2016, in which Purdue Pharma agreed to pay Practice Fusion approximately $1 million. The CDSS program went live on July 6, 2016, containing three separate alerts. The first alert encouraged the providers to document a pain score for the patient. The second alert provided a "Brief Pain Inventory" with the patient's previous pain scores within the previous three months. Then the provider would go through a list of questions on the severity and impact of patient's pain, summarizing the patient's current pain as "worst," "on average," and "least" in the previous 24 hours. The third alert would appear if the patient reported a pain score of four or higher twice within four months, or the patient was diagnosed with chronic pain after completing the Brief Pain Inventory.

In the third alert, the CDSS utilized a drop-down menu of options for pain management.

[1] "Practice Fusion Lands A Whopping $70M To Bring A Big Data Cure To The Healthcare Crisis," TechCrunch, September 2013.

[2] All details and quotes in this section, unless otherwise specified, are from Exhibit C "Statement of Facts" of Case 2:20-cr-00011-wks, United States Attorney's Office for the District of Vermont.

The list was sourced from a 2016 New England Journal of Medicine article titled "Opioid Abuse in Chronic Pain — Misconceptions and Mitigation Strategies" (Volkow and McLellan, 2016), although the authors explicitly warn that the intention is not to provide a clinical instruction in chronic pain management. in addition to the non-opioid analgesics and other pain-management options listed in the article, Practice Fusion added the opioid therapy without considering the patient's condition. In particular, ERO products are listed for patients with less than severe or non-chronic pain, and even in cases where the pain could be treated with non-ERO options.

On December 14, 2016, Practice Fusion conducted a presentation at Purdue Pharma's headquarter. The meeting revealed that through November 30, 2016, the CDSS had alerted during 21 million visits involving 7.5 million patients, generating a general shift toward EROs particularly from immediate-release opioids. The CDSS alerted more than 230,000,000 times from July 6, 2016 to Spring 2019, in which it was suspended right before Purdue Pharma filed for bankruptcy.

2.2 EHR Adoption

Our data construction starts with a granular record of eligible physicians adopting different EHR products from April 2011 to September 2017, downloaded from the HealthIT.gov website. In 2009, Congress passed and President Obama signed the Health Information Technology for Economic and Clinical Health (HITECH) Act, encouraging the adoption of health IT to improve the quality and efficiency of care. The Act created incentive programs that offer physicians financial benefits for the meaningful use of EHRs. To start, the CMS and the Office of the National Coordinator for Health Information Technology (ONC) have established standards for EHRs to become certified. Practice Fusion's EHRs were certified in 2011. Eligible professionals (EPs), including most U.S. physicians, can adopt certified EHR technology and attest to fulfilling meaningful use criteria to receive incentive payments. Each physician must achieve three stages with different criteria to receive full payments.[3]

We focus on the Medicare EHR Incentive Program administered by the CMS from 2011

[3]These stages include data capturing and sharing (stage 1), advanced clinical process (stage 2), and improved outcomes (stage 3).

to 2016, providing a maximum of $44,000 in total payments across years.[4] A participating EP must demonstrate meaningful use every year to receive payments, and the CMS publicly releases a complete record of attestation information, enabling researchers to track each EP's EHR products over time. The HealthIT.gov streamlines and disambiguates the raw data into a single "EHR Products Used for Meaningful Use Attestation" file. In each attestation year, this file lists all the certified EHR products attested by an EP (identified by the NPI). For each record, it supplements the provider zip code, specialty, vendor name, product version and product setting. The initial sample contains over 1.8 million records for almost 360,000 NPIs using 1,232 EHR products owned by 730 vendors.

We impose the following major restrictions on the initial sample. First, certified EHR products are utilized in various application categories, including ambulatory, cardiology imaging, financial decision support, human resources and IS security, etc. All Practice Fusion products are ambulatory EHRs. To maintain comparability, we drop all the EHR products without any ambulatory applications. There are over 1 million adoption records of the remaining 190 products by roughly 210,000 physicians. Second, we focus on the regions (3-digit zip code areas) with non-zero physicians adopting Practice Fusion products to alleviate the confounding effects due to heterogeneous healthcare conditions across regions. Lastly, we drop physicians using four vendors with other misbehaviors in the Incentive Program, including Athenahealth, Nextgen Healthcare, Modernizing Medicine (ModMed), and Greenway Health. Their violations include missing the required functionality to become certified and paying physicians to falsely attest.[5] None of these vendors promoted opioids or received kickbacks from opioid drug makers. The rationale is that physician practices may be distorted by these EHR products as well, questioning their validity as clean control groups. The final sample has roughly 27% (57,138) of the remaining physicians following the first step. To sum up, these are physicians participating in the Medicare EHR Incentive Program and having meaningfully used EHRs with ambulatory functions. They either adopt

[4] Alternatively, EPs can participate in the Medicaid EHR Incentive Program run by every state. However, there is no centralized disclosure platform for Medicaid programs. But they can only participate in one program.

[5] The only exception is that ModMed solicited and received kickbacks from Miraca Life Sciences Inc. (Miraca) in exchange for recommending Miraca's pathology lab services. However, we cannot perform the analysis based on this case since we do not have utilization data of Miraca's services. Our results are also robust to keeping physicians using these four vendors.

Practice Fusion products (treatment group), or are comparable neighbors of Practice Fusion users (control group).

2.3 Opioid Prescription

We then link the above sample to the "Medicare Part D Prescribers by Provider" dataset by the CMS, containing annual prescription information by individual physicians under the Medicare Part D Prescription Drug Program from 2013 to 2021. Relevant to our study, it provides the total number of claims, costs, and days of supply of opioid drugs each year. Besides, it provides the same information for long-acting opioids. These are drugs formulated to release more gradually into the bloodstream with a longer duration of analgesic action. Purdue Pharma's dominant ERO products, e.g., OxyContin, are all long-acting. We back out the short-acting (SA) opioid claim amount for a placebo test. As these drugs are not promoted by Practice Fusion, the treated physicians should not exhibit significant changes in SA opioid prescriptions throughout the sample.

We include a few patient characteristics as control variables. These are the average age of the beneficiaries, percent of male beneficiaries, percent of African American beneficiaries, percent of Hispanic beneficiaries, percent of beneficiaries qualified to receive both Medicare and Medicaid benefits, and the average risk score of beneficiaries. The last variable is a hierarchical condition category developed by the CMS based health-influencing factors.[6]

2.4 Purdue Pharma Detailing

Other than manipulating the CDSS, Purdue Pharma's main marketing strategy is direct detailing to physicians. Indeed, since Purdue Pharma introduced OxyContin in 1996, it aggressively engaged in detailing to promote the product, leading to the decades-long opioid crisis (Van Zee, 2009; Alpert et al., 2022). We obtain Purdue Pharma detailing information from the CMS Open Payments database. This database collects and reports financial relationships between drug and medical device companies and providers since 2013, involving

[6]These factors include the beneficiary's age, sex, eligibility for Medicaid, initial reason for Medicare qualification, residence in an institution such as a long-term care facility, and the diagnoses assigned to the beneficiary in inpatient, outpatient and office-based settings during a base year.

detailing payments such as research, meals, travel, gifts or speaking fees. For physicians in our sample, we collect whether they receive payments from Purdue Pharma in a given year. Section 4.2 compares the long-term prescription behavior impacts of this explicit marketing strategy with the implicit manipulation of the decision support system.

3 Empirical Design

Our main sample comprises annual observations of physician EHR usage and opioid prescriptions from 2013 to 2021. This sample spans three phases: pre-treatment (2013 – 2015), treatment (2016 – 2018), and post-treatment (2019 – 2021). Our main specification is essentially a simple DID design:

$$Y_{i,k,t} = \alpha + \beta PF_i \times Post_t + \gamma' X_{i,t} + \delta_i + \mu_{k,t} + \varepsilon_{i,k,t}. \tag{1}$$

In the above equation, $Y_{i,k,t}$ characterizes the opioid prescription of physician i from area k in year t. PF_i indicates whether physician i adopted Practice Fusion products in the Medicare EHR Incentive Program. $Post_t$ is one if year t is greater or equal to 2016, and zero otherwise. $X_{i,t}$ indicates the set of control variables capturing patient characteristics, detailed in Section 2.3. We include two-way fixed effects (FEs) — the physician FE, δ_i, and the area-year FE, $\mu_{k,t}$ — which will absorb the standalone effects of PF_i and $Post_t$. This specification effectively compares the opioid prescription of two otherwise similar physicians in the same area and year, where one is exposed to the Practice Fusion manipulation, and the other is not.

To disentangle the short-term and long-term effects, we estimate equation (1) separately in the following two samples. The first sample is truncated at 2018, including only the pre-treatment and treatment phases, and demonstrating the direct effects of CDSS manipulation. The second sample drops the yearly observations from 2016 to 2018, consisting of the pre-treatment and post-treatment phases, and exhibiting the persistent impacts on human capital. Table 1 Panel A summarizes the main variables in the whole sample (all three phases). 6,244 unique physicians, i.e. more than 10% of the sample, belong to the

treatment group. On average, each physician prescribes 13.35 annual LA claims, supplying 377 days of usage and costing around $2,600 for all the beneficiaries. LA claims account for 4.26% of all the opioid claims. Since our sample focuses on the Medicare Part D prescription, the average beneficiary in our sample is above 70 years old and more risky (139%) than the national average (normalized as 1). Unconditionally, each physician has a 4% chance of receiving promotional payments from Purdue Pharma.

Our identification relies on the parallel trend assumption, which requires that the treatment and control groups exhibit similar tendencies in LA opioid prescriptions without Practice Fusion's impacts. This assumption likely holds in our setting for the following reasons. First, the self-selection concern is mitigated by the timing difference. Treated physicians participated in the EHR incentive program and adopted Practice Fusion's products in 2011. The actual manipulation occurred in 2016. It is almost implausible that physicians were able to predict the manipulation or intentionally adopted a product generating biased recommendations. Second, there exists an alternative explanation that the treatment group had a stronger preference for LA opioids. They insisted on prescribing Oxycontin after Purdue Pharma suffered from litigation and reputation losses in recent years, while the control group was more sensitive to the negative news. Contrary to this hypothesis, Table 1 Panel B shows that, if anything, the treatment group were more conservative in LA opioids. Compared to the control group, they supplied significantly 52 fewer days of drugs, and their prescription rate was also 1% lower. Lastly, we will further visualize the coefficient dynamics, demonstrating no significant trends before the shock. Besides, we perform a placebo test that shows the null effects on SA opioids, i.e., products not recommended in Practice Fusion's CDSS.

4 Results

4.1 Opioid Prescription

We start by estimating equation (1) in the initial sample from 2013 to 2018 in Table 2 to investigate whether the treatment group significantly prescribes more LA opioids when Practice Fusion manipulated the CDSS. This analysis provides a first-stage result to con-

firm the existence of direct impacts in the short term. One may argue that as professionals with expertise, affected physicians will use their discretion to reject the manipulated recommendation and refrain from opioid prescription in unnecessary situations. Instead, Table 2 indicates that physicians utilizing Practice Fusion EHRs experienced a significant 7.6% increase in annual LA claims compared to the control group, equivalent to approximately one additional claim each year, from 2016 to 2018. Meanwhile, their total LA opioid costs surged by 19.2%, accompanied by a 14.3% increase in days of supply. A ballpark estimate of the total additional costs for Medicare Part D in our sample is approximately $3.14 million every year.[7] This magnitude is close to Practice Fusion's own prediction, between $8,458,232 and $11,277,643, in opioid revenue driven by the CDSS.

Table 2 also includes a placebo test to validate our identification assumption. As we explained earlier, Purdue Pharma's leading products are extended-release opioids utilized for chronic pain management, suggesting the manipulation should have minimal effects on SA opioids. Column (4) verifies that the fraction of LA claims to total opioid claims increases by 0.29 percentage points relative to the control group, corresponding to a 7% jump from the sample average (4.26%). Lastly, column (5) shows that affected physicians do not exhibit any significant differences in SA claims after the shock, confirming that the changes were concentrated in LA opioids. The negative coefficient, though insignificant, is in line with Practice Fusion's claim that the CDSS alert generated a shift from SA opioids to ERO products.

Next, we estimate equation (1) in the sample consisting of both pre-treatment and post-treatment phases. If the distortion on opioid prescription was only transient, then affected physicians would not exhibit any significant changes in LA claims from 2019 to 2021. However, Table 3 demonstrates that affected physicians continue to prescribe more LA opioid claims even after the CDSS stopped making recommendation after the Spring of 2019. The economic magnitudes are even larger. For example, column (1) implies a 12.4% increase in LA claims, which is over 160% of the corresponding coefficient in Table 2. Rises in LA opioid costs for affected physicians become 32.1%, implying a $5.2 million aggregated additional

[7]The calculation is based on the sample average annual cost $2,595, multiplied by the coefficient and the number of physicians in the treatment group (6,244).

payments by Medicare Part D every year in our sample. This result implies that, through learning by doing, the short-term artificial intelligence manipulation will generate a persistent long-term impacts on human capital. Affected physicians form the habit of opioid usage and ultimately maintain the prescription tendency, regardless of the CDSS recommendation.

Tables 2 and 3 estimate the average treatment effects in different phases. To visualize the coefficient dynamics in the whole sample, we estimate and plot the coefficients β_c from the following equation in Figure 1:

$$Y_{i,k,t} = \alpha + \sum_{\substack{c\in[2013,2021] \\ c\neq 2015}} \beta_c PF_i \times \mathbf{I}(t=c) + \gamma' X_{i,t} + \delta_i + \mu_{k,t} + \varepsilon_{i,k,t}. \tag{2}$$

$\mathbf{I}(t=c)$ indicates whether year t is calendar year c. β_c estimates the difference between the treatment and control group at calendar time c. Following the convention, we drop the year 2015, the year before the treatment, making it as the reference year in the above equation. There are two main takeaways from the figure. First, there exist no significant differences between the two groups in the pre-treatment phase (before the red vertical line), supporting the parallel-trend assumption. Second, while Practice Fusion manipulated the CDSS (between the two vertical lines), affected physicians gradually increased LA opioid usage over the years. While this momentum stopped after the suspension of biased recommendation in 2019, the treatment group maintained a constantly high frequency of prescription. These dynamics explain the magnitude differences in Tables 2 and 3.

Lastly, we supplement additional evidence on the long-term impacts from the perspective of real costs on beneficiaries. Starting from 2020, Medicare covers Opioid Treatment Programs (OTPs) that provide medication-assisted treatment for people diagnosed with an opioid use disorder (OUD) in Medicare Part B. Through the "Medicare Physician & Other Practitioners" file, we can aggregate the area-level Medicare utilization of OTPs based on the HCPCS codes and use it as a proxy for local opioid addictions among the elderly group. We prefer this measure over other public data, such as mortality due to drug overdose, for two reasons. First, although we cannot track the actual OUD conditions of the specific beneficiaries who received LA opioids from the treatment group, the aggregated utilization is a fuzzy but consistent measure from the same Medicare group. Second, overdose mortality is

an extreme situation that occurs with a lower frequency, leaving us with smaller variation in the data.

Because the coverage began in recent years and did not span the pre-treatment phase, we acknowledge the limitation of this analysis as correlational evidence and interpret the results with caution. In Table 4, our focal regressor is the number of unique physicians that adopted Practice Fusion in the Medicare Incentive Program in each 3-digit zip code area. The outcome variables include the number of beneficiaries, number of services, bill charges and actual Medicare payments for OTPs. Control variables include local demographic characteristics such as logged population, logged income per capital, male percentage, African American percentage, Hispanic percentage, average age, unemployment rate, and health insurance coverage percentage. We find that Practice Fusion usage among providers indeed correlates with long-term OUDs among the Medicare beneficiaries. For example, a 1% increase in the number of Practice Fusion adopters is associated with 0.48% more Medicare beneficiaries utilizing OTPs.

4.2 Comparison to Detailing

In this section, we demonstrate that human capital distortion exists only when physicians are unaware of the biased recommendation. To establish, we delve into an alternative group of physicians that receive explicit in-kind payments from Purdue Pharma and evaluate their LA opioid prescription behaviors over time. As shown in Figure A.1, Purdue Pharma detailing activities peaked at 2014, where almost 4,000 physicians in our sample received promotion payments. The number of recipients steadily decreased since then, and only one physician maintained a financial relationship with the company in 2019. There exist two potential reasons for this trend. First, the CMS Open Payments system was implemented online in 2014, retrospectively revealing all the in-kind payments from August 2013. Physicians might worry about their public images being damaged by disclosed payments. Second, since Purdue Pharma's settlement with Kentucky in 2015, additional states sued it at a rapid speed. By January 2019, 36 states were suing Purdue Pharma. The potential litigation risk might make the company more cautious with promotion.

Recall that our treatment group, i.e., Practice Fusion adopters, continue increasing LA

opioid prescription even in the background of Purdue Pharma's negative publicity in this window. In Table 5 Panel A, we replicate the analysis from Table 3 analysis by comparing the prescription behaviors of payment recipients in the post-treatment phases to those in the pre-treatment phase. $Detail_i$ indicates whether physician i has received in-kind payments before 2019. Note that Purdue Pharma effectively suspended detailing activities after filing for bankruptcy in 2019. Therefore, Table 5 allows us to evaluate the possibility that detailing recipients will form a similar habit and maintain the high frequency of opioid prescription without financial incentives anymore.

However, former payment recipients substantially reduce LA opioid prescriptions after 2019. Their LA claim amount on average drops by 41% compared to the pre-treatment period before 2015. Indeed, Figure 2 shows that their prescription frequency steadily decreased since 2015, consistent with the trend that Purdue Pharma detailing became less popular in the same window. This result highlights that even though financial incentives can also contribute to transient jumps in LA opioid usage, they cannot foster a long-term preference. Physicians may simply decide on prescriptions as a function of financial benefits and even worry about litigation and reputation when lawsuits started to target Purdue Pharma. Therefore, we argue that a key feature of artificial intelligence on human capital is the unconsciousness of behavior changes. Professionals are less likely to endogenously correct their working habits without focusing on explicit incentives.

Figure A.2 shows that the geographic distribution of the two physician groups is highly uncorrelated ($\rho = -0.19$), further confirming our argument that Practice Fusion adopters did not have other incentives to prescribe LA opioids. As a robustness check, Table 5 Panel B shows that our main results in Table 3 are consistent when controlling for the physician receiving detailing payments. From Purdue Pharma's perspective, manipulating Practice Fusion's CDSS was a substitute of detailing for promotion, partially offsetting its revenue losses.

4.3 Heterogeneity

Which types of affected physicians are less sensitive to the implicit impacts from the EHR alert? In this section, we explore several heterogeneity tests based on ex-ante characteristics

to better understand the interaction between human and artificial intelligence. We start with focusing on the experience of affected physicians for two general reasons. First, during Practice Fusion's communication with Purdue Pharma, they believed the CDSS would target "opioid naive" users and utilize their limited knowledge on the potential addiction risks. Second, senior physicians are more likely to make decisions by their own expertise and have a lower trust for the artificial intelligence.

In Table 6, we interact the treatment variable with the pre-shock level of LA opioid claims for each physician in 2015. The coefficients of all interaction terms are negative with substantial statistical significance. To understand the economic magnitudes, Panel A column (1) implies that a typical treated physician with a pre-existing median-level LA prescription amount (0) will have a treatment effect of 22.6% in the short-term. This impact will reduce to zero if that physician had around 3 LA claims before the shock. Other than the aforementioned two general reasons, it is also possible that if physicians had previously prescribed opioids, they might be more attentive to the coverage of ERO's risks and reluctant to accept opioids as a treatment option.

Table 7 studies a broader seniority without focusing on the specific interactions with opioids. We define seniority as the number of years since medical school graduation. We obtain similar effects with Table 6: affected physicians with a higher seniority are less sensitive to the CDSS manipulation. However, the economic magnitudes are much smaller, and the coefficients are not significant for $LARate$, highlighting the importance of opioid-specific expertise. Admittedly, our measure of seniority could also be a proxy for physician age. Existing literature has documented that older physicians are less likely to utilize EHRs (Kane and Labianca, 2011; Venkatesh et al., 2011), which can be an alternative channel explaining the results. These two tables overall suggest that human capital formation among physicians can be highly path-dependent. Doctors at the early-stage of their careers are more sensitive to external manipulations and can repeat mistakes into long-term habits.

Tables 8 and 9 assess the importance of demand-side patient characteristics by examining beneficiary risk score and age. The rationale is that physicians may exercise additional caution and invest more effort in considering the side effects of extended-release opioids when facing patients with higher risk profiles. Both tables support this hypothesis as almost

all interaction terms are significantly negative, except for columns (4) in Table 9. In terms of magnitudes, column (1) of Table 8 Panel B suggests that increasing the ex-ante patient risk score from the 25th percentile (1.08) to the 75th percentile (1.55) will decrease the long-term impacts on LA claims by 6.3%. Beneficiary age has a slightly greater impact, as changing the ex-ante patient age from the 25th percentile (70.09) to the 75th percentile (74.40) can reduce the same outcome by 7.3%.

5 Conclusion

This paper sheds light on the lasting impacts of AI interactions on human capital in the context of opioid prescriptions by physicians assisted with clinical decision support systems. A biased pain alert promoting the sales of extended-release opioids resulted in a significant increase in opioid claims by affected physicians during the treatment window (2016–2018). This effect persisted even after the removal of the biased alert, indicating a long-term distortion in human capital development. Notably, this distortion was not observed in physicians receiving explicit payments from Purdue Pharma, highlighting the unique role of AI in shaping human behavior beyond traditional financial incentives.

The study contributes to understanding the dynamics of AI interactions on professionals' long-term beliefs and habits. The unconscious nature of AI biases makes them undetectable even by professionals with expertise. Instead, physicians may perceive the recommendations as optimally generated from real data and update more favorable beliefs regarding the benefit-risk profile of opioids over time. Our study serves as a caveat for the potential risks of AI manipulation and the need for greater awareness and accountability in AI-driven decision-making processes. In light of AI's expanding influence across diverse sectors, managing its effects on human capital is essential for ensuring responsible and ethical technological integration.

References

Acemoglu, Daron, 2021, Harms of ai, Technical report, National Bureau of Economic Research.

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo, 2022, Artificial intelligence and jobs: evidence from online vacancies, *Journal of Labor Economics* 40, S293–S340.

Acemoglu, Daron, and Pascual Restrepo, 2018, The race between man and machine: Implications of technology for growth, factor shares, and employment, *American economic review* 108, 1488–1542.

Acemoglu, Daron, and Pascual Restrepo, 2020, Robots and jobs: Evidence from us labor markets, *Journal of political economy* 128, 2188–2244.

Acemoglu, Daron, and Pascual Restrepo, 2022, Tasks, automation, and the rise in us wage inequality, *Econometrica* 90, 1973–2016.

Agarwal, Nikhil, Alex Moehring, Pranav Rajpurkar, and Tobias Salz, 2023, Combining human expertise with artificial intelligence: Experimental evidence from radiology, Technical report, National Bureau of Economic Research.

Agrawal, Ajay, Joshua S Gans, and Avi Goldfarb, 2019, Artificial intelligence: the ambiguous labor market impact of automating prediction, *Journal of Economic Perspectives* 33, 31–50.

Alpert, Abby, William N Evans, Ethan MJ Lieber, and David Powell, 2022, Origins of the opioid crisis and its enduring impacts, *The Quarterly Journal of Economics* 137, 1139–1179.

Angelova, Victoria, Will S Dobbie, and Crystal Yang, 2023, Algorithmic recommendations and human discretion, Technical report, National Bureau of Economic Research.

Autor, David, and Anna Salomons, 2018, Is automation labor-displacing? productivity growth, employment, and the labor share, Technical report, National Bureau of Economic Research.

Bresnahan, Timothy F, and Manuel Trajtenberg, 1995, General purpose technologies 'engines of growth'?, *Journal of econometrics* 65, 83–108.

Bronsoler, Ari, Joseph Doyle, and John Van Reenen, 2022, The impact of health information and communication technology on clinical quality, productivity, and workers, *Annual Review of Economics* 14, 23–46.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock, 2018, What can machines learn and what does it mean for occupations and the economy?, in *AEA papers and proceedings*, volume 108, 43–47, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Cao, Sean, Wei Jiang, Junbo L Wang, and Baozhong Yang, 2021, From man vs. machine to man+ machine: The art and ai of stock analyses, Technical report, National Bureau of Economic Research.

Coleman, Braiden, Kenneth Merkley, and Joseph Pacelli, 2022, Human versus machine: A comparison of robo-analyst and traditional research analyst investment recommendations, *The Accounting Review* 97, 221–244.

Cowgill, Bo, and Catherine E Tucker, 2019, Economics, fairness and algorithmic bias, *preparation for: Journal of Economic Perspectives* .

Doyle, Joseph, Sarah Abraham, Laura Feeney, Sarah Reimer, and Amy Finkelstein, 2019, Clinical decision support for high-cost imaging: a randomized clinical trial, *PLoS One* 14, e0213373.

Dranove, David, and Craig Garthwaite, 2022, Artificial intelligence, the evolution of the healthcare value chain, and the future of the physician, Technical report, National Bureau of Economic Research.

D'Acunto, Francesco, Nagpurnanand Prabhala, and Alberto G Rossi, 2019, The promises and pitfalls of robo-advising, *The Review of Financial Studies* 32, 1983–2020.

Felten, Edward W, Manav Raj, and Robert Seamans, 2018, A method to link advances in artificial intelligence to occupational abilities, in *AEA Papers and Proceedings*, volume 108, 54–57, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Fuster, Andreas, Paul Goldsmith-Pinkham, Tarun Ramadorai, and Ansgar Walther, 2022, Predictably unequal? the effects of machine learning on credit markets, *The Journal of Finance* 77, 5–47.

Ganju, Kartik K, Hilal Atasoy, Jeffery McCullough, and Brad Greenwood, 2020, The role of decision support systems in attenuating racial biases in healthcare delivery, *Management science* 66, 5171–5181.

Grennan, Jillian, and Roni Michaely, 2020, Artificial intelligence and high-skilled work: Evidence from analysts, *Swiss Finance Institute Research Paper* .

Hydari, Muhammad Zia, Rahul Telang, and William M Marella, 2019, Saving patient ryan—can advanced electronic medical records make patient care safer?, *Management Science* 65, 2041–2059.

Kane, Gerald C, and Giuseppe Labianca, 2011, Is avoidance in health-care groups: A multilevel investigation, *Information Systems Research* 22, 504–522.

Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan, 2018, Human decisions and machine predictions, *The quarterly journal of economics* 133, 237–293.

Mullainathan, Sendhil, and Ziad Obermeyer, 2022, Diagnosing physician error: A machine learning approach to low-value health care, *The Quarterly Journal of Economics* 137, 679–727.

Rossi, Alberto G, and Stephen P Utkus, 2020, Who benefits from robo-advising? evidence from machine learning, *Evidence from Machine Learning (March 10, 2020)* .

Van Zee, Art, 2009, The promotion and marketing of oxycontin: commercial triumph, public health tragedy, *American Journal of Public Health* 99, 221–227.

Venkatesh, Viswanath, Xiaojun Zhang, and Tracy A Sykes, 2011, "doctors do too little technology": A longitudinal field study of an electronic healthcare system implementation, *Information Systems Research* 22, 523–546.

Volkow, Nora D, and A Thomas McLellan, 2016, Opioid abuse in chronic pain—misconceptions and mitigation strategies, *New England Journal of Medicine* 374, 1253–1263.

Webb, Michael, 2019, The impact of artificial intelligence on the labor market, *Available at SSRN 3482150* .

Figure 1: Treatment Dynamics of Practice Fusion's Manipulation

This figure plots the treatment dynamics with respect to LA opioid prescription. We plot the β_cs estimated from equation (2). x-axis indicates the calendar years. The base year is 2015 (indicated by the red vertical line), that is, the year prior to the shock. The green vertical indicates the suspension of Practice Fusion manipulation (2019). 95% confidence intervals are indicated by the solid lines.

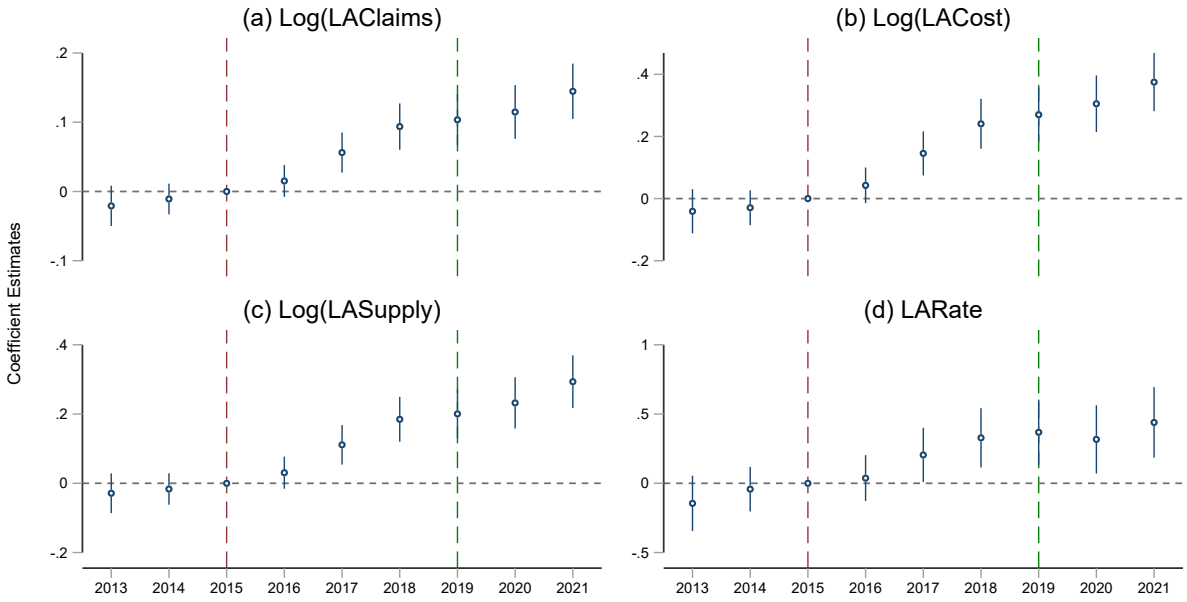


Figure 2: Treatment Dynamics of Purdue Pharma Detailing

This figure plots the treatment dynamics with respect to LA opioid prescription due to Purdue Pharma detailing. As before, we estimate β_cs in equation (2), except that we replace PF_i with $Detail_i$. The base year is 2015 (indicated by the red vertical line), that is, the year prior to the shock. The green vertical indicates the suspension of in-kind payments (2019). 95% confidence intervals are indicated by the solid lines.

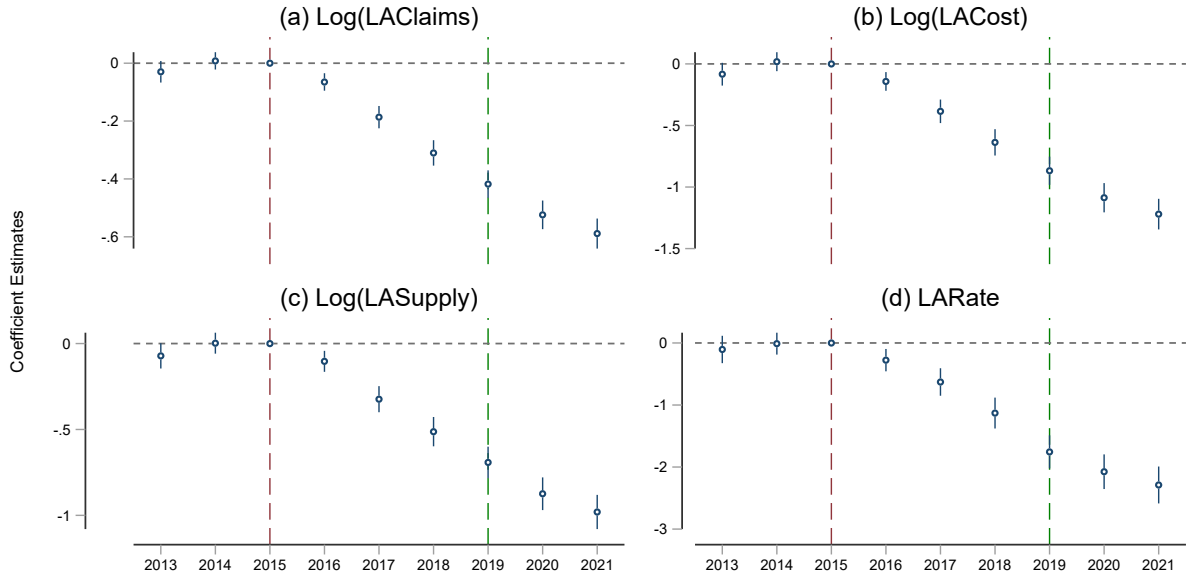


Table 1: Summary Statistics

This table provides the summary statistics of the main variables in the paper. Panel A summarizes the annual observations of physicians from 2013 to 2021. PF_i is one if physician i adopts Practice Fusion's EHR, and zero otherwise. $Post_t$ is one if year t is greater or equal to 2016, and zero otherwise. $LAClaims_{i,t}$ is the number of long-acting opioid claims by physician i in year t. $LACost_{i,t}$ is the dollar amount of long-acting opioid prescription costs by physician i in year t. $LASupply_{i,t}$ is the days of supply for long-acting opioids by physician i in year t. $LARate_{i,t}$ is the percentage of long-acting opioid claims out of all opioid claims by physician i in year t. $AvgAge_{i,t}$ is the average age of beneficiaries visiting physician i in year t. $MalePct_{i,t}$ is the fraction of male beneficiaries visiting physician i in year t. $BlackPct_{i,t}$ is the fraction of African American beneficiaries visiting physician i in year t. $HispanicPct_{i,t}$ is the fraction of Hispanic beneficiaries visiting physician i in year t. $AvgRisk_{i,t}$ is the average risk score of beneficiaries visiting physician i in year t. $DualPct_{i,t}$ is the fraction of beneficiaries qualified to receive both Medicare and Medicaid benefits. $Detailing_{i,t}$ is one physician i receives in-kind payments from Purdue Pharma in year t, and zero otherwise. Panel B compares the treatment and control groups in 2015, the year before the treatment. t-statistics are calculated based on standardard errors clustered at the 3-digit zip code level.

Panel A: Full Sample

Variable	N	Mean	Std	p25	Median	p75
$PF \times Post$	427,104	0.03	0.18	0.00	0.00	0.00
$LAClaims$	427,104	13.35	32.93	0.00	0.00	12.00
$LACost$	427,104	2,595.40	7,628.17	0.00	0.00	747.94
$LASupply$	427,104	376.97	919.79	0.00	0.00	360.00
$LARate$	427,104	4.26	8.95	0.00	0.00	4.52
$AvgAge$	427,104	71.93	3.93	70.09	72.42	74.40
$MalePct$	427,104	0.40	0.13	0.32	0.41	0.48
$BlackPct$	427,104	0.11	0.18	0.00	0.03	0.14
$HispanicPct$	427,104	0.07	0.15	0.00	0.00	0.07
$AvgRisk$	427,104	1.39	0.51	1.08	1.26	1.55
$DualPct$	427,104	0.22	0.18	0.08	0.18	0.32
$Detailing$	427,104	0.04	0.19	0.00	0.00	0.00

Panel B: Ex-ante Difference in 2015

	Control Obs.	*Treat Obs.*	*Control Mean*	*Treat Mean*	*Diff.*	*t-stat*	*p-value*
$LAClaims$	44,727	5,602	16.57	15.46	1.11	1.28	0.20
$LACost$	44,727	5,602	3,169.53	2,940.71	228.82	1.30	0.20
$LASupply$	44,727	5,602	471.73	419.62	52.11	2.17	0.03
$LARate$	44,727	5,602	5.16	4.21	0.95	4.36	0.00
$Detailing$	44,727	5,602	0.07	0.07	0.00	0.67	0.50

Table 2: The Short-term Impacts on Opioid Prescription

This table provides the results on the short-term impacts of Practice Fusion's CDSS on opioid prescription. The sample consists of annual observations at the physician level from 2013 to 2018, i.e. the pre-treatment and treatment phases. PF_i is one if physician i adopts Practice Fusion's EHR, and zero otherwise. $Post_t$ is one if year t is greater or equal to 2016, and zero otherwise. $Log(LAClaims)_{i,t}$ is the logarithm of one plus $LAClaims_{i,t}$. $Log(LACost)_{i,t}$ is the logarithm of one plus $LACost_{i,t}$. $Log(LASupply)_{i,t}$ is the logarithm of one plus $LASupply_{i,t}$. $LARate_{i,t}$ is the percentage of long-acting opioid claims out of all opioid claims by physician i in year t. $Log(SAClaims)_{i,t}$ is the logarithm of one plus number of short-acting opioid claims. Control variables include $AvgAge$, $MalePct$, $BlackPct$, $HispanicPct$, $DualPct$, and $AvgRisk$. Physician fixed effects and area-year fixed effects are included. Standard errors are clustered at the physician level, and t-statistics are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

	(1) *Log(LAClaims)*	(2) *Log(LACost)*	(3) *Log(LASupply)*	(4) *LARate*	(5) *Log(SAClaims)*
$PF \times Post$	0.076***	0.192***	0.143***	0.288***	−0.018
	(5.892)	(6.377)	(5.868)	(3.618)	(−1.259)
Controls	Y	Y	Y	Y	Y
Area×Year FEs	Y	Y	Y	Y	Y
Physician FEs	Y	Y	Y	Y	Y
N	339,454	339,454	339,454	339,454	339,454
R^2	0.78	0.74	0.74	0.68	0.86

Table 3: The Long-term Impacts on Opioid Prescription

This table provides the results on the long-term impacts of Practice Fusion's CDSS on opioid prescription. The sample consists of annual observations at the physician level from 2013 to 2015, and 2019 to 2021, i.e. the pre-treatment and post-treatment phases. PF_i is one if physician i adopts Practice Fusion's EHR, and zero otherwise. $Post_t$ is one if year t is greater or equal to 2016, and zero otherwise. $Log(LAClaims)_{i,t}$ is the logarithm of one plus $LAClaims_{i,t}$. $Log(LACost)_{i,t}$ is the logarithm of one plus $LACost_{i,t}$. $Log(LASupply)_{i,t}$ is the logarithm of one plus $LASupply_{i,t}$. $LARate_{i,t}$ is the percentage of long-acting opioid claims out of all opioid claims by physician i in year t. Control variables include $AvgAge$, $MalePct$, $BlackPct$, $HispanicPct$, $DualPct$, and $AvgRisk$. Physician fixed effects and area-year fixed effects are included. Standard errors are clustered at the physician level, and t-statistics are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

	(1) *Log(LAClaims)*	(2) *Log(LACost)*	(3) *Log(LASupply)*	(4) *LARate*
$PF \times Post$	0.124***	0.321***	0.241***	0.399***
	(6.926)	(7.779)	(7.212)	(3.679)
Controls	Y	Y	Y	Y
Area×Year FEs	Y	Y	Y	Y
Physician FEs	Y	Y	Y	Y
N	281,310	281,310	281,310	281,310
R^2	0.73	0.69	0.69	0.62

Table 4: Long-term Opioid Therapy Utilization

This table provides the cross-sectional evidence on the long-term costs for Medicare beneficiaries due to Practice Fusion adoption. The sample consists of annual observations at the 3-digit zip code level in 2020 and 2021. $Log(NumPF)_k$ is number of physicians having adopted Practice Fusion in area k. $OTPBenes_{k,t}$, $OTPSvcs_{k,t}$, $OTPCharge_{k,t}$, and $OTPPay_{k,t}$ are the number of beneficiaries, number of services, bill charges and actual Medicare payments for Opioid Treatment Programs in area k and year t, respectively. Control variables include local demographic characteristics such as logged population, logged income per capital, male percentage, African American percentage, Hispanic percentage, average age, unemployment rate and health insurance coverage percentage. Year fixed effects are included. Standard errors are clustered at the area level, and t-statistics are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

	(1) *Log(OTPBenes)*	(2) *Log(OTPSvcs)*	(3) *Log(OTPCharge)*	(4) *Log(OTPPay)*
$Log(NumPF)$	0.475***	0.763***	1.269***	1.245***
	(4.584)	(4.451)	(4.459)	(4.424)
Controls	Y	Y	Y	Y
Year FEs	Y	Y	Y	Y
N	1,777	1,777	1,777	1,777
R^2	0.28	0.28	0.27	0.27

Table 5: Comparison of Long-term Effects: Explicit Detailing v.s. Implicit Manipulation

This table provides the results on the long-term impacts of Purdue Pharma's detailing on opioid prescription. The sample consists of annual observations at the physician level from 2013 to 2015, and 2019 to 2021, i.e. the pre-treatment and post-treatment phases. $Detail_i$ is one if physician i had received in-kind payments from Purdue Pharma by 2019, and zero otherwise. PF_i is one if physician i adopts Practice Fusion's EHR, and zero otherwise. $Post_t$ is one if year t is greater or equal to 2016, and zero otherwise. $Log(LAClaims)_{i,t}$ is the logarithm of one plus $LAClaims_{i,t}$. $Log(LACost)_{i,t}$ is the logarithm of one plus $LACost_{i,t}$. $Log(LASupply)_{i,t}$ is the logarithm of one plus $LASupply_{i,t}$. $LARate_{i,t}$ is the percentage of long-acting opioid claims out of all opioid claims by physician i in year t. Control variables include $AvgAge$, $MalePct$, $BlackPct$, $HispanicPct$, $DualPct$, and $AvgRisk$. Physician fixed effects and area-year fixed effects are included. Standard errors are clustered at the physician level, and t-statistics are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Effects of Detailing

	(1) *Log(LAClaims)*	(2) *Log(LACost)*	(3) *Log(LASupply)*	(4) *LARate*
$Detail \times Post$	−0.407***	−0.773***	−0.638***	−2.132***
	(−5.800)	(−4.901)	(−5.088)	(−5.316)
Controls	Y	Y	Y	Y
Area×Year FEs	Y	Y	Y	Y
Physician FEs	Y	Y	Y	Y
N	272,892	272,892	272,892	272,892
R^2	0.71	0.68	0.68	0.61

Panel B: Comparison of Effects

	(1) *Log(LAClaims)*	(2) *Log(LACost)*	(3) *Log(LASupply)*	(4) *LARate*
$PF \times Post$	0.184***	0.440***	0.341***	0.759***
	(10.394)	(10.614)	(10.093)	(7.004)
$Detail \times Post$	−0.565***	−1.152***	−0.931***	−2.785***
	(−7.911)	(−7.152)	(−7.269)	(−6.795)
Controls	Y	Y	Y	Y
Area×Year FEs	Y	Y	Y	Y
Physician FEs	Y	Y	Y	Y
N	272,892	272,892	272,892	272,892
R^2	0.71	0.68	0.68	0.61

Table 6: Heterogeneous Effects of Impacts: Long-acting Opioid Experience

This table provides the results on the heterogeneous impacts due to long-acting opioid prescription experience before the shock. Panel A estimates the effects in the short-term sample (2013 – 2018) and Panel B in the long-term sample (2013 – 2015 & 2019 – 2021). PF_i is one if physician i adopts Practice Fusion's EHR, and zero otherwise. $Post_t$ is one if year t is greater or equal to 2016, and zero otherwise. $Log(LAClaims)_i^{2015}$ is the logarithm of one plus long-acting opioid claims for physicians i in year 2015, i.e. the pre-shock level of LA opioid usage. $Log(LAClaims)_{i,t}$ is the logarithm of one plus $LAClaims_{i,t}$. $Log(LACost)_{i,t}$ is the logarithm of one plus $LACost_{i,t}$. $Log(LASupply)_{i,t}$ is the logarithm of one plus $LASupply_{i,t}$. $LARate_{i,t}$ is the percentage of long-acting opioid claims out of all opioid claims by physician i in year t. Control variables include $AvgAge$, $MalePct$, $BlackPct$, $HispanicPct$, $DualPct$, and $AvgRisk$. Physician fixed effects and area-year fixed effects are included. Standard errors are clustered at the physician level, and t-statistics are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Short-term Effects

	(1) *Log(LAClaims)*	(2) *Log(LACost)*	(3) *Log(LASupply)*	(4) *LARate*
$PF \times Post$	0.226***	0.504***	0.400***	1.092***
	(21.770)	(19.688)	(19.426)	(17.078)
$PF \times Post$	−0.162***	−0.335***	−0.277***	−0.865***
$\times Log(LAClaims)^{2015}$	(−18.212)	(−17.169)	(−17.793)	(−14.852)
Controls	Y	Y	Y	Y
Area×Year FEs	Y	Y	Y	Y
Physician FEs	Y	Y	Y	Y
N	338,944	338,944	338,944	338,944
R^2	0.78	0.74	0.75	0.68

Panel B: Long-term Effects

	(1) *Log(LAClaims)*	(2) *Log(LACost)*	(3) *Log(LASupply)*	(4) *LARate*
$PF \times Post$	0.437***	0.963***	0.774***	2.047***
	(33.447)	(30.560)	(30.304)	(24.779)
$PF \times Post$	−0.330***	−0.678***	−0.563***	−1.739***
$\times Log(LAClaims)^{2015}$	(−25.980)	(−23.722)	(−24.541)	(−22.988)
Controls	Y	Y	Y	Y
Area×Year FEs	Y	Y	Y	Y
Physician FEs	Y	Y	Y	Y
N	280,758	280,758	280,758	280,758
R^2	0.73	0.70	0.70	0.62

Table 7: Heterogeneous Effects of Impacts: Physician Seniority

This table provides the results on the heterogeneous impacts due to physician seniority. Panel A estimates the effects in the short-term sample (2013 – 2018) and Panel B in the long-term sample (2013 – 2015 & 2019 – 2021). PF_i is one if physician i adopts Practice Fusion's EHR, and zero otherwise. $Post_t$ is one if year t is greater or equal to 2016, and zero otherwise. $Log(Seniority)_i^{2015}$ is the logarithm of one plus number of years since graduation for physicians i in year 2015, i.e. the pre-shock level of seniority. $Log(LAClaims)_{i,t}$ is the logarithm of one plus $LAClaims_{i,t}$. $Log(LACost)_{i,t}$ is the logarithm of one plus $LACost_{i,t}$. $Log(LASupply)_{i,t}$ is the logarithm of one plus $LASupply_{i,t}$. $LARate_{i,t}$ is the percentage of long-acting opioid claims out of all opioid claims by physician i in year t. Control variables include $AvgAge$, $MalePct$, $BlackPct$, $HispanicPct$, $DualPct$, and $AvgRisk$. Physician fixed effects and area-year fixed effects are included. Standard errors are clustered at the physician level, and t-statistics are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Short-term Effects

	(1) *Log(LAClaims)*	(2) *Log(LACost)*	(3) *Log(LASupply)*	(4) *LARate*
$PF \times Post$	0.269***	0.623***	0.489***	0.207
	(3.347)	(3.384)	(3.282)	(0.433)
$PF \times Post$	−0.058**	−0.131**	−0.105**	0.032
$\times Log(Seniority)^{2015}$	(−2.383)	(−2.336)	(−2.318)	(0.221)
Controls	Y	Y	Y	Y
Area×Year FEs	Y	Y	Y	Y
Physician FEs	Y	Y	Y	Y
N	324,047	324,047	324,047	324,047
R^2	0.78	0.75	0.75	0.68

Panel B: Long-term Effects

	(1) *Log(LAClaims)*	(2) *Log(LACost)*	(3) *Log(LASupply)*	(4) *LARate*
$PF \times Post$	0.485***	1.156***	0.921***	0.642
	(4.352)	(4.570)	(4.487)	(0.936)
$PF \times Post$	−0.110***	−0.256***	−0.208***	−0.070
$\times Log(Seniority)^{2015}$	(−3.215)	(−3.286)	(−3.297)	(−0.334)
Controls	Y	Y	Y	Y
Area×Year FEs	Y	Y	Y	Y
Physician FEs	Y	Y	Y	Y
N	268,060	268,060	268,060	268,060
R^2	0.73	0.70	0.70	0.62

Table 8: Heterogeneous Effects of Impacts: Patient Risk

This table provides the results on the heterogeneous impacts due to patient risks. Panel A estimates the effects in the short-term sample (2013 – 2018) and Panel B in the long-term sample (2013 – 2015 & 2019 – 2021). PF_i is one if physician i adopts Practice Fusion's EHR, and zero otherwise. $Post_t$ is one if year t is greater or equal to 2016, and zero otherwise. $Log(AvgRisk)_i^{2015}$ is the logarithm of one plus average beneficiary risk score for physicians i in year 2015, i.e. the pre-shock level of patient risk. $Log(LAClaims)_{i,t}$ is the logarithm of one plus $LAClaims_{i,t}$. $Log(LACost)_{i,t}$ is the logarithm of one plus $LACost_{i,t}$. $Log(LASupply)_{i,t}$ is the logarithm of one plus $LASupply_{i,t}$. $LARate_{i,t}$ is the percentage of long-acting opioid claims out of all opioid claims by physician i in year t. Control variables include $AvgAge$, $MalePct$, $BlackPct$, $HispanicPct$, $DualPct$, and $AvgRisk$. Physician fixed effects and area-year fixed effects are included. Standard errors are clustered at the physician level, and t-statistics are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Short-term Effects

	(1) *Log(LAClaims)*	(2) *Log(LACost)*	(3) *Log(LASupply)*	(4) *LARate*
$PF \times Post$	0.220***	0.420***	0.342***	0.820**
	(4.388)	(3.513)	(3.588)	(2.507)
$PF \times Post$	−0.156***	−0.246**	−0.216**	−0.576*
$\times Log(AvgRisk)^{2015}$	(−2.955)	(−1.977)	(−2.174)	(−1.656)
Controls	Y	Y	Y	Y
Area×Year FEs	Y	Y	Y	Y
Physician FEs	Y	Y	Y	Y
N	338,944	338,944	338,944	338,944
R^2	0.78	0.74	0.75	0.68

Panel B: Long-term Effects

	(1) *Log(LAClaims)*	(2) *Log(LACost)*	(3) *Log(LASupply)*	(4) *LARate*
$PF \times Post$	0.410***	0.720***	0.614***	1.580***
	(5.699)	(4.335)	(4.559)	(3.589)
$PF \times Post$	−0.309***	−0.431**	−0.403***	−1.275***
$\times Log(AvgRisk)^{2015}$	(−4.065)	(−2.488)	(−2.863)	(−2.811)
Controls	Y	Y	Y	Y
Area×Year FEs	Y	Y	Y	Y
Physician FEs	Y	Y	Y	Y
N	280,758	280,758	280,758	280,758
R^2	0.73	0.69	0.70	0.62

Table 9: Heterogeneous Effects of Impacts: Patient Age

This table provides the results on the heterogeneous impacts due to patient age. Panel A estimates the effects in the short-term sample (2013 – 2018) and Panel B in the long-term sample (2013 – 2015 & 2019 – 2021). PF_i is one if physician i adopts Practice Fusion's EHR, and zero otherwise. $Post_t$ is one if year t is greater or equal to 2016, and zero otherwise. $Log(AvgAge)_i^{2015}$ is the logarithm of one plus average beneficiary age for physicians i in year 2015, i.e. the pre-shock level of patient age. $Log(LAClaims)_{i,t}$ is the logarithm of one plus $LAClaims_{i,t}$. $Log(LACost)_{i,t}$ is the logarithm of one plus $LACost_{i,t}$. $Log(LASupply)_{i,t}$ is the logarithm of one plus $LASupply_{i,t}$. $LARate_{i,t}$ is the percentage of long-acting opioid claims out of all opioid claims by physician i in year t. Control variables include $AvgAge$, $MalePct$, $BlackPct$, $HispanicPct$, $DualPct$, and $AvgRisk$. Physician fixed effects and area-year fixed effects are included. Standard errors are clustered at the physician level, and t-statistics are reported in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Short-term Effects

	(1) *Log(LAClaims)*	(2) *Log(LACost)*	(3) *Log(LASupply)*	(4) *LARate*
$PF \times Post$	2.387***	5.192***	4.542***	1.446
	(2.850)	(2.845)	(3.084)	(0.237)
$PF \times Post$	−0.540***	−1.167***	−1.027***	−0.270
$\times Log(AvgAge)^{2015}$	(−2.760)	(−2.738)	(−2.984)	(−0.190)
Controls	Y	Y	Y	Y
Area×Year FEs	Y	Y	Y	Y
Physician FEs	Y	Y	Y	Y
N	338,944	338,944	338,944	338,944
R^2	0.78	0.74	0.75	0.68

Panel B: Long-term Effects

	(1) *Log(LAClaims)*	(2) *Log(LACost)*	(3) *Log(LASupply)*	(4) *LARate*
$PF \times Post$	5.455***	11.746***	9.915***	5.454
	(4.559)	(4.511)	(4.698)	(0.620)
$PF \times Post$	−1.245***	−2.667***	−2.258***	−1.180
$\times Log(AvgAge)^{2015}$	(−4.455)	(−4.386)	(−4.581)	(−0.576)
Controls	Y	Y	Y	Y
Area×Year FEs	Y	Y	Y	Y
Physician FEs	Y	Y	Y	Y
N	280,758	280,758	280,758	280,758
R^2	0.73	0.69	0.70	0.62

Appendix

A Appendix Figures

Figure A.1: Number of Recipients of Purdue Pharma Detailing

This figure plots the yearly number of recipients of Purdue Pharma detailing in our sample. We stopped at 2019 since in-kind payments to physicians were effectively suspended after Purdue Pharma filed for bankruptcy in 2019.

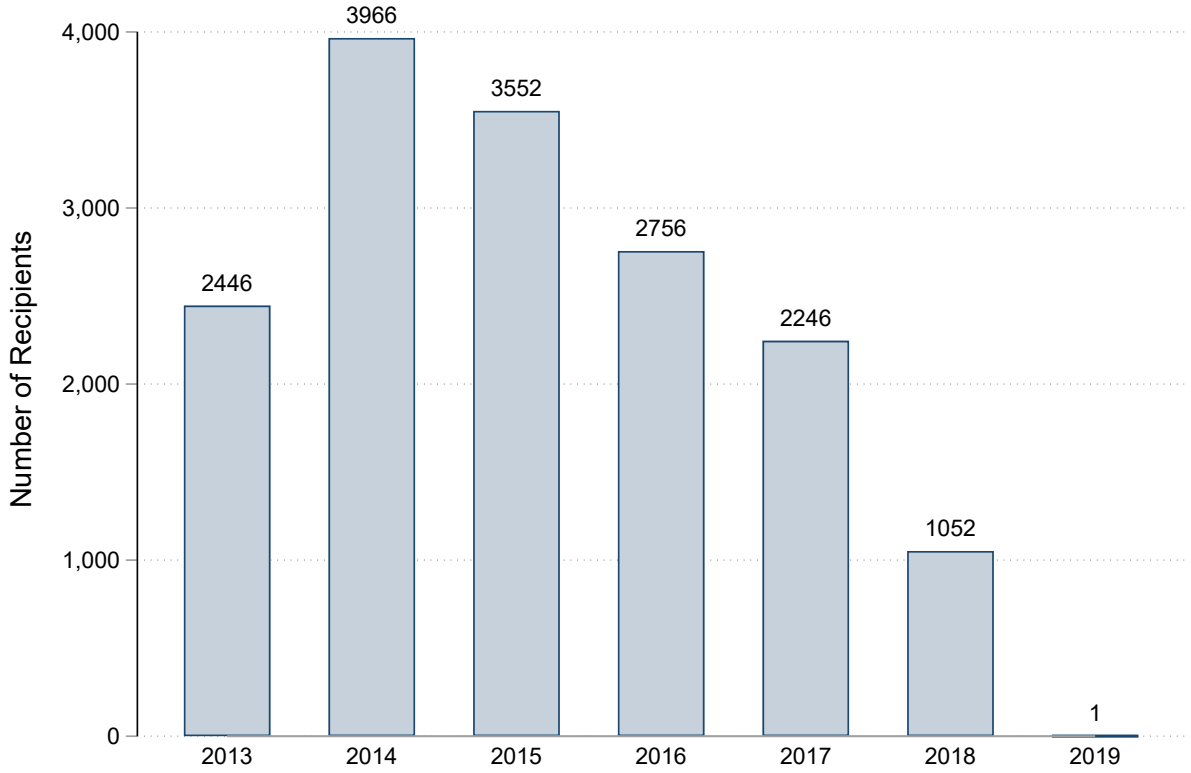
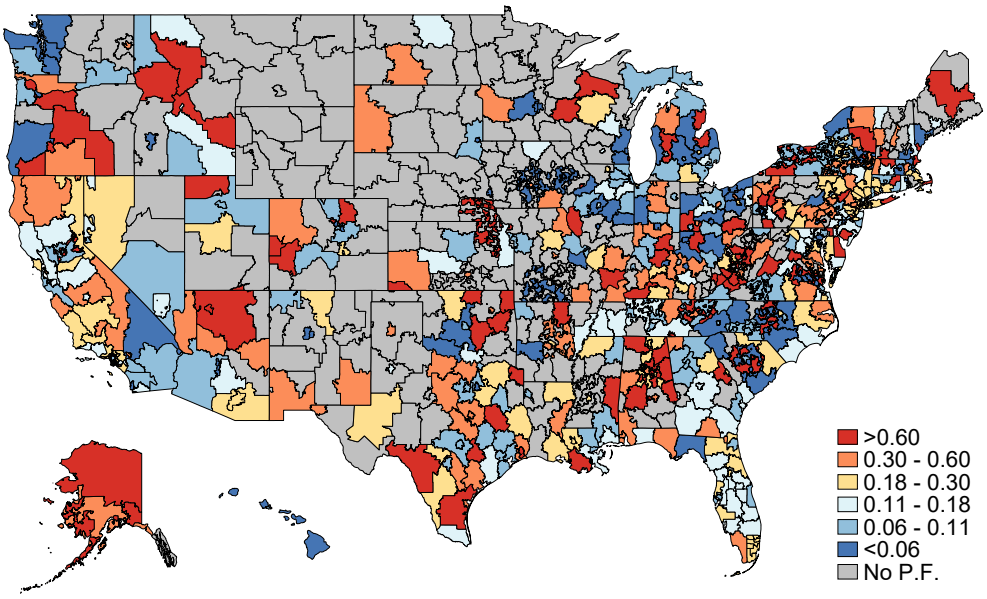


Figure A.2: Number of Recipients of Purdue Pharma Detailing

This figure plots the geogrphaci distributioon of Practice Fusion adopters and detailing recipients in our sample. Each area is a 3-digit zip code area. In each area, we calculate the fraction of Practice Fusion adopters and detailing recipients over the total number of physicians in that area from our sample. Areas without Practice Fusion adopters are marked in gray and dropped from the sample in our analysis.

Panel A: Distribution of Practice Fusion Adopters



Panel B: Distribution of Detailing Recipients

