Aligning Incentives or Gaming the System? The Impact of Insurer-Physician Acquisitions *

Link to latest version

Nick Grasley and Mariana Guido†

November 25, 2025

Abstract

Vertical integration between insurers and primary care practices is increasingly common in the US healthcare market, raising concerns among policymakers. We use medical claims data from the Colorado All Payer Claims Database to study a 2017 insurer-physician acquisition, focusing on how aligning incentives within the integrated firm affects physician behavior. We find two main impacts. First, in Medicare Advantage, where insurer revenue is tied to patients' diagnoses, we leverage a patient-level event study design to estimate how integration increases diagnosis-based payments to the insurer by $998 to $1,805 per patient per year, without a corresponding increase in treatment. Second, as the integrated firm bears the cost of specialist care, we estimate a referral choice model to evaluate whether the integrated practice steers referrals towards more cost-effective specialists. We find that, for patients in the Commercial segment, the acquired practice steers referrals and saves approximately $300 per inpatient referral. Our results reveal a crucial trade-off: integration can increase taxpayer costs in one market while generating efficiency gains in another. Furthermore, we show that contracting at non-acquired practices can replicate the diagnostic coding effect. This substitutability between integration and contracting suggests that policies targeting only mergers may be ineffective at addressing strategic coding behavior.

*We thank Liran Einav, Matthew Gentzkow, and José Ignacio Cuesta for their mentorship. We thank Claudia Allende, Laurence Baker, Sebastián Bauer, Matthew Brown, Dante Domenella, Harsh Gupta, Helen Kissel, Neale Mahoney, Maria Polyakova, Adam Rosenberg, Maya Rossin-Slater, Gregory Rosston, Helena Roy, Adrienne Sabety, Jason Weitze, and seminar participants at Stanford University for valuable comments and suggestions. We thank the Center for Improving Value in Healthcare, especially Lucia Sanders, Hannah Witting, and Kristin Paulson, for assistance in providing and interpreting the data. Data for this project were accessed using the Stanford Center for Population Health Sciences Data Core. The PHS Data Core is supported by a National Institutes of Health National Center for Advancing Translational Science Clinical and Translational Science Award (UL1TR003142) and from Internal Stanford funding. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH. We gratefully acknowledge financial support from the George P. Shultz Dissertation Support Fund and the Program in Regulatory Policy at the Stanford Institute for Economic Policy Research.

†Department of Economics, Stanford University

1 Introduction

Vertical integration between physicians and insurers has become increasingly common over the last decade, raising concerns for policymakers and regulators (DOJ, 2024; FTC, 2024; WSJ et al., 2025). Although historically vertical integration in healthcare was led by hospital acquisitions of physicians (Handel and Ho, 2021; Cooper et al., 2025), currently, the largest employers of physicians in the United States are insurers, with the largest employing approximately 10% of all physicians (KFF, 2024). Despite the growing prevalence of these acquisitions, we have limited evidence quantifying their impacts.

The welfare impacts of vertical integration are both complex and theoretically ambiguous (Riordan, 2008; Crawford et al., 2018). Integration may impact competitive forces outside the firm – e.g., through reduced double marginalization and foreclosure (Williamson, 1971; Hart and Tirole, 1990) – as well as align incentives within the integrated firm (Williamson, 1985; Grossman and Hart, 1986). The latter is of particular interest in the healthcare market, where insurers and physicians face conflicting incentives over the provision of care (Frandsen and Rebitzer, 2019). Whereas insurers would like for physicians to fully internalize the financial impacts of their care decisions, the typical physician payment structure incentivizes the provision of *more* care (Dranove, 1988).

This paper examines how vertical integration aligns incentives between insurers and physicians and impacts physician behavior in the context of a 2017 insurer-physician acquisition in Colorado using rich medical claims data. We consider two main margins of changes in physician behavior following integration, with implications for both the Medicare Advantage and Commercial segments. First, in the Medicare Advantage segment, where diagnosis-based payments are a key source of revenue for insurers (Geruso and Layton, 2020), we conduct an event study design to show that the acquired practice provides more diagnoses, delivering increased payments. Second, as the integrated firm bears the cost of specialist care across market segments, we evaluate whether the acquired practice steers referrals towards more cost-effective specialists. To do this, we estimate a model of physician referrals, where we find steering among Commercial beneficiaries. Combining these findings highlights a tradeoff across market segments in our setting: whereas in Medicare Advantage integration increases costs for taxpayers, in the Commercial segment integration increases efficiency.

To guide our empirical analysis of the effects of vertical integration, we develop a conceptual framework that illustrates how the insurer can pass through its incentives to the physician. Our framework highlights two main changes in physicians' behavior. First, the physician increases their provision of diagnoses. This occurs in the Medicare Advantage segment, a private insurer alternative to Traditional Medicare for individuals aged 65 and older, where the government pays the insurer as a function of physicians' diagnostic provision. Therefore, the integrated insurer-physician, by increasing its diagnostic provision, increases government payments and taxpayer costs. Second, the physician reduces referral costs, generating savings for the insurer in specialist care. This occurs in both the Medicare Advantage and Commercial (e.g., employer-sponsored health insurance) segments, where the physician controls access to specialist care through referrals and the insurer must reimburse specialists for this care. Lastly, our conceptual framework highlights how contracting may achieve the same form of incentive alignment as vertical integration.

We empirically investigate these impacts of insurer-primary care physician vertical integration in the context of one such instance in Colorado in 2017. We leverage an extract from the Colorado All Payer Claims Database (CO APCD), covering patients enrolled across all private insurance plans from 2015 to 2019, including both the Medicare Advantage and Commercial market segments. These data and acquisition are uniquely well-suited for our purposes. Not only do the data provide rich coverage of patients' care and outcomes at the acquired practice before and after integration, but they also include detailed information from non-integrated practices and insurers. The data include detailed line-item information for each medical encounter, such as procedure codes and payment amounts, patient and provider information, and, crucially, an identifier for the acquiring insurer. Furthermore, the acquisition we study is representative of the insurer's acquisitions in this time period.

We begin our empirical analysis by focusing on the acquired practice's diagnostic provision and exploring descriptive patterns in the market. Two key patterns emerge. First, the acquisition coincides with the widespread adoption of pay-per-patient contracts with the acquiring insurer, both at the acquired practice and at non-acquired primary care practices. These contracts specify a physician payment arrangement that can replicate the incentives from vertical integration through diagnosis-based payments and cost-sharing for the primary care physician (Ho and Pakes, 2014a). Second, we find an upward trend in diagnoses for Medicare Advantage beneficiaries of the acquiring insurer. This trend holds for patients at both the acquired practice and non-acquired practices, coinciding with the acquisition and the widespread adoption of pay-per-patient contracts.

These descriptive findings highlight the key challenges in identifying the impact of incentive alignment through vertical integration on diagnosis-based payments to the insurer in Medicare Advantage. Not only is the acquisition endogenous, but in this setting, acquisition can be interpreted as a means to facilitate contract adoption. This highlights that other practices that adopt pay-per-patient contracts in the absence of acquisition may be fundamentally different from the acquired practice, making them ill-suited as a comparison group. We overcome this challenge by leveraging our cross-insurer patient-level panel data. We compare beneficiaries of the acquiring insurer to those of other insurers within the acquired practice. In this way, we identify how the incentive alignment enabled by acquisition impacts physicians' differential provision of diagnoses for beneficiaries of the acquiring insurer. The identifying assumption underlying our research design is that, on average, the health condition of patients of the acquired practice across insurers evolves in the same way. Thus, differences in their diagnoses over time result from changes in physicians' diagnostic behavior following integration.[1]

[1] Note that, as we cannot observe patients' true underlying health status, we do not observe whether, after the acquisition, patients are over-diagnosed, or if they were under diagnosed before.

We estimate the impact of vertical integration using an event study design with individual and time fixed effects, following Callaway and Sant'Anna (2021). We find that, as illustrated by our conceptual framework, the acquired practice provides increased diagnoses to the Medicare Advantage beneficiaries of the acquiring insurer. Indeed, our estimates imply increased payments to the insurer of $998 to $1,805 per patient per year, approximately a 30% increase of baseline payments per year. We compare these estimates to the impact of pay-per-patient contract adoption at non-acquired practices and find that these are statistically indistinguishable. We argue that these two results do not imply that contracting is a perfect substitute for acquisition, as practices face differential selection into acquisition and contract adoption. Furthermore, we find substantial heterogeneity in the impact of contracting on diagnostic behavior across practices, suggesting that selection may matter for the magnitude of these impacts.

We then leverage our conceptual framework to interpret the implications of our findings. We start by noting that the increase in payments to the insurer, while increasing insurer profits, also increases the burden on taxpayers. Next, we consider the possibility that increased diagnoses may be accompanied by increased treatment, which could be beneficial for patients' health. We evaluate this both through our framework and empirically. First, we expand our conceptual framework to include treatment decisions alongside diagnostic ones. Our framework predicts that, even if treatment increases, it will lag behind the increase in diagnoses, as financial incentives for the latter are larger. Second, we empirically confirm this prediction. We find that the acquired practice increases their provision of diagnostic procedures, e.g., blood tests, approximately five times as much as their provision of treatment procedures, e.g., dialysis. This divergence suggests that the increase in diagnoses is unlikely to be driven by greater patient care needs.

Next, we shift our focus to evaluating changes in physicians' referral patterns, where the acquired practice internalizes the cost of specialist care as illustrated by our conceptual framework. We start by constructing a sample of referrals for both inpatient and outpatient specialist services. Then we explore descriptive patterns in specialist care around integration. First, we note that the scope for cost savings from changes in referral behavior is inherently constrained by the level and variation in specialist prices. Note that prices in Medicare Advantage are anchored to prices in Traditional Medicare (Trish et al., 2017), whereas Commercial prices are not. Consistent with this, we find that prices in the Commercial segment exhibit higher mean and variance than in Medicare Advantage, indicating greater scope for savings in the Commercial segment. Second, although we find an increase in the number of referrals among Medicare Advantage beneficiaries of the acquiring insurer, we do not find any evidence of the same among Commercial beneficiaries.

To further explore changes in physicians' referral patterns, we leverage our detailed price data to focus on whether the acquired practice indeed steers its patients towards more cost-effective specialists. The key challenge to evaluating whether the acquired practice steers referrals towards more cost-effective specialists is that, while we observe realized costs, physicians make decisions based on expected costs. We overcome this by estimating a referral model in which physicians and patients jointly choose which specialist to visit by trading off expected costs, quality, and distance. We first construct specialists' expected costs for each service by leveraging all referrals to each specialist, and then we perform both a leave-one-out correction and hierarchical Bayes shrinkage. We leverage this cost index to estimate physicians' price sensitivity for beneficiaries of the acquiring insurer and other insurers, allowing it to vary over time and across beneficiary types. These estimates capture how the acquired practice's price sensitivity for referrals of the acquiring insurer's beneficiaries changes after integration.

We find that the acquired practice becomes more price sensitive when referring Commercial beneficiaries of the acquiring insurer. Among inpatient referrals, this increased price sensitivity translates into savings of approximately $300 per referral, or 18% of the average cost per referral. Although we find heterogeneous estimates across different categories of outpatient specialist care, our estimates imply average savings of approximately $26 per referral, or 50% of the average cost per referral. However, consistent with the reduced scope for savings in Medicare Advantage, we find no statistically significant savings in this segment, either at the acquired practice or at practices that adopted pay-per-patient contracts. Insofar as integration is better suited to overcome informational asymmetries, it may be more effective at achieving the desired referral steering than contracting. However, we do not observe any contract adoption in the Commercial segment, the only segment in which we find evidence of referral steering, so we cannot directly speak to this hypothesis.

To evaluate the overall implications of vertical integration in this market, we combine our findings across physician behaviors and market segments. Our findings in the Medicare Advantage segment are consistent with the expected increase in diagnoses and diagnosis-based payments for beneficiaries of the acquiring insurer. However, we find no evidence of cost savings in referrals in Medicare Advantage. In contrast, we find evidence of cost savings for Commercial beneficiaries of the acquiring insurer. Ultimately, the welfare implications of this integration are ambiguous, involving a trade-off between increased costs for taxpayers in Medicare Advantage and potential efficiency gains for private payers in the Commercial market, which may be passed through to consumers. Furthermore, the full welfare implications of integration will also depend on other impacts of vertical integration, for instance, on competition and patient health.

Lastly, we highlight the policy implications of our findings. While much of the recent policy discussion has focused on antitrust enforcement against vertical mergers in healthcare (FTC, 2024; Dafny, 2014), our study demonstrates the importance of jointly considering integration and contracting. Our results show that, in the context of diagnostic coding, pay-per-patient contracts can act as effective substitutes for ownership in aligning incentives. This finding is consistent with both the theoretical literature on the internal organization of the firm (Grossman and Hart, 1986) and its application in the context of healthcare and physician responses to financial incentives (Gaynor et al., 2004). This substitutability implies that blocking integration may be an ineffective policy lever if the firm can achieve a similar outcome by adopting a pay-per-patient contract instead. Furthermore, although contracts are adopted only in one market segment, integration affects both Medicare Advantage and Commercial beneficiaries, with efficiency gains in the latter segment through referral steering. Therefore, a comprehensive regulatory approach must address incentive alignment in both its integrated and contractual forms and its implications across market segments.

This paper is closely related to three strands of the literature. First, there is an extensive literature on vertical integration, both theoretically (Perry, 2008; Rey and Tirole, 2007; Riordan, 2008; Bresnahan and Levin, 2012) and empirically (Hortaçsu and Syverson, 2007; Atalay et al., 2014; Crawford et al., 2018; Luco and Marshall, 2020). In particular, a part of this literature focuses on vertical integration in the healthcare setting. There exists a large literature exploring the impacts of integration between hospitals and physicians (Baker et al., 2016; Brot-Goldberg and de Vaan, 2018; Capps et al., 2018; Cutler et al., 2020; Koch et al., 2021; Cooper et al., 2025) and between insurers and hospitals (Johnson et al., 2017; Diebel, 2018; Park et al., 2023; Cuesta et al., 2024). They find that this form of integration can lead to price increases and the foreclosure of non-integrated market participants. A comparatively small portion of the literature considers integration between insurers and primary care physicians. Closest to our setting, Cho (2025) examines the integration of insurers and healthcare providers, including primary care physicians, and provides evidence on its impact on referral steering. We contribute to this literature by considering the heterogeneity of these impacts across market segments and by further studying the impact of integration on coding in Medicare Advantage.

Second, we contribute to the literature on regulatory gaming, which examines the manipulability of regulation in the healthcare setting (Dafny and Dranove, 2009; Eliason et al., 2019; Gupta and Sacarny, 2025) as well as in other economic settings such as credit rating (Griffin and Tang, 2011) and pollution auditing (Duflo et al., 2013). More specifically, we focus on the context of Medicare Advantage and risk-adjustment. This literature examines healthcare spending and revenues in Medicare Advantage (Brown et al., 2014; Curto et al., 2019) as well as coding incentives (Decarolis, 2015; Einav et al., 2016; Fang and Gong, 2017; Decarolis et al., 2020). Most closely related to our paper, Geruso and Layton (2020) finds that Medicare Advantage plans generate higher risk scores than traditional Medicare, and this increase in coding is larger among provider-owned plans. We build on this finding by documenting the causal impact of vertical integration on risk-adjustment.

Lastly, our results on how integration aligns incentives between insurers and physicians, as well as their interaction with pay-per-patient contracts build on the literature on the internal organization of firms (Grossman and Hart, 1986) and its implications in the health care industry (Gaynor et al., 2004; Frandsen and Rebitzer, 2014). Specifically, this paper is related to the literature on how physicians respond to financial incentives (Clemens and Gottlieb, 2014; Ho and Pakes, 2014a,b; Einav et al., 2018; Gupta, 2021) which examines how physician remuneration impacts health care supply and expenditures.

The rest of the paper proceeds as follows. In Section 2, we describe the institutional setting and the data we use. In Section 3, we write down a simple model elucidating on how integration may mitigate agency problems between insurers and physicians. Section 4 presents our estimates of the impact of the acquisition on diagnosis-based payments to the insurer in the Medicare Advantage segment. In Section 5 we estimate the impact of the acquisition on the cost of specialist care. Lastly, we discuss the overall implications of the acquisition across segments and how they generalize outside our setting in Section 6 before concluding in Section 7.

2 Background and Data

2.1 Background

In the U.S. healthcare market, individuals pay a premium to enroll in private insurance plans in the Commercial segment – mostly obtained through employer-sponsored health insurance or the individual market – or in the Medicare Advantage segment (MA) – a private insurance option that provides an alternative to Traditional Medicare.[2] Insurers negotiate with healthcare providers over network inclusion and prices, determining which physicians their beneficiaries can receive care from and at what price. To access non-emergency care, beneficiaries typically start by visiting their primary care physician, who manages their overall health and controls referrals to downstream specialists and hospitals.

In MA, insurers receive a Risk Adjustment payment per beneficiary from the Center for Medicare and Medicaid Services (CMS) as a function of their beneficiaries' health status, meant to compensate insurers for enrolling costlier beneficiaries. CMS determines this payment by multiplying patients' risk scores by a base rate. Patients' risk scores, which are supposed to reflect their expected healthcare costs, are calculated using patients' prior-year diagnoses and demographic information such as age and gender. The base rate is determined at the county level and are the result of a competitive bidding process. MA plans' revenues consist primarily of this fixed risk-adjusted payment per beneficiary and they must pay for the cost of care incurred by their beneficiaries.[3] These features incentivize plans to provide high-value care, keeping their beneficiaries healthy and managing their healthcare costs (Alliance, 2017), while also maximizing their beneficiaries' risk adjustment payments.

In both the MA and Commercial segments, insurers' healthcare costs consist of payments to physicians for their care provision. There are two main payment arrangements between insurers and physicians: fee-for-service (FFS) and pay-per-patient (or capitation) contracts.[4,5] Under FFS,

[2]Traditional Medicare, also known as fee-for-service Medicare, is the federal health insurance program, administered by The Centers for Medicare and Medicaid Services (CMS), meant to cover care for individuals 65 or older, as well as those with some disabilities.

[3]As in other insurance markets, beneficiaries wishing to enroll in a plan must pay an insurance premium. However, in practice, premiums are often zero in Medicare Advantage.

[4]In practice, payment arrangements lie somewhere in the spectrum between these two extremes.

[5]These types of contracts are commonly referred to as capitation contracts and typically apply to patients in health

insurers and providers negotiate over a payment rate per procedure. Then, physicians are paid per service provided, which may incentivize over-provision of care (Dranove, 1988).

In contrast, under pay-per-patient contracts, insurers and providers negotiate over a fixed risk-adjusted per-patient rate meant to cover all of the patient's care. Broadly, these contracts specify a budget meant to cover all of the patient's expenditures, either solely at the physician practice or, in the case of "global-budget" contracts, across all of their care utilization, and may also include incentive payments for improved quality of care. In this way, physicians are made to bear some share of the cost of their treatment decisions (Brot-Goldberg and de Vaan, 2018). Furthermore, these budgets are set as a function of patients' underlying health status, as reflected in their risk scores. Given this structure, these contracts are seen as a powerful tool for insurers to manage costs by passing through their incentives to physicians.

An alternative mechanism for passing through these financial incentives has become increasingly prominent over the past decade: vertical integration between insurers and physician practices. From 2019 to 2023, the number of physician practices owned by insurers and private equity investors doubled and insurers are now the largest (individual) employers of physicians, surpassing hospital systems (PAI, 2024).[6] Indeed, recent testimony to Congress (Damberg, 2023) highlights how insurers have become increasingly vertically integrated over the past decade, allowing them to "capture the revenues created along the [health care] production path".

2.2 Setting and Data

We study a single insurer's acquisition of a primary care practice (PCP) in Denver, Colorado, in 2017. This event provides an ideal setting for our analysis for two reasons. First, the acquisition is representative of the insurer's broader acquisition strategy during this period (see Appendix Table F.1). Second, its location in Colorado allows us to leverage rich, panel data from the state's All-Payer Claims Database to measure its impacts.

Our primary dataset is an extract from the Colorado All Payer Claims Database (CO APCD).[7] Under state law, all insurers are required to submit processed medical claims to this database, which is administered by the Center for Improving Value in Health Care (CIVHC). The dataset is extensive, covering Commercial and Medicare Advantage Claims from 2015 to 2019. It contains 41 million claims for 1.5 million unique patients and 58,000 medical providers, representing $16.9 billion in total spending. The data include 48 insurance plans, of which 7 belong to the acquiring

maintenance organizations (HMOs), which require a designated primary care physician for each patient. Note that whereas the term "capitation" is used in various healthcare settings to describe how physicians are paid, we focus on contracts between private insurers and physicians, using slightly different nomenclature to distinguish between them.

[6] Measures of consolidation have grouped insurers and private equity together under "Corporations". We focus on insurer acquisitions, noting that private equity acquisitions have distinct ownership incentives compared to those of insurers.

[7] These data, and comparable data from other states' APCDs, have been used in numerous recent research (Liebman and Panhans, 2021; Ghili et al., 2023; Cho, 2025).

insurer.[8,9]

The claims data provide us with detailed information on enrollees, medical providers, and every reimbursed medical encounter between them. For each claim, we observe the medical procedures performed (CPT codes), patient diagnoses (ICD codes), and payments from both the insurer and the patient.[10] The data also indicates whether a claim was paid under a pay-per-patient contract.[11,12] Crucially, the CO APCD contains a unique patient identifier that allows us to track individuals over time, even as they switch between different insurance plans. The data also includes a distinct identifier for the acquiring insurer's plans.

We supplement the data with a few external sources. We use data from PitchBook to observe firms structure and access information on the acquisition. To accurately track physician practices over time, we use historical snapshots of the National Plan and Provider Enumeration System (NPPES), made available by the NBER. Finally, we leverage zipcode-level income information from the 2016 Internal Revenue Service Statistics of Income to construct income measures.

Our analysis requires a dataset that links patients to their primary care practices and tracks their outcomes over time. To construct this sample, we perform three main tasks. First, we group individual physicians into their respective practices leveraging physicians' reported "business mailing address" and accounting for physicians changing practices over time.[13] Second, we identify each patient's primary care practice in the pre- and post-acquisition periods (2015-2016 and 2016-2019, respectively) as the practice at which they received the most primary care each period. Lastly, we construct patient risk scores using the 2019 CMS-HCC software, which, roughly, computes a weighted sum of patients' diagnoses in a year, accounting for demographics (age and gender). We provide more details on these procedures in Appendix A.

2.3 Sample Descriptives

From our original extract of the CO APCD (1.5 million patients), we first restrict our sample to the 1.3 million with any medical claims. We further restrict the sample to the 1.19 million patients who are ever assigned a primary care provider, our *Full Sample*, of which 743 thousand are assigned a primary care provider both before and after the acquisition. These patients are assigned to 2,687 distinct primary care practices. Practice size is fairly skewed: while the average practice is small

[8] Providers include medical professionals as well as the organizations in which they practice, e.g., hospitals and physician practices.

[9] We exclude claims from certain insurers (and their beneficiaries) for whom the data are not representative of their patients' care, as well as secondary and tertiary claims.

[10] Our data separately identify plan-paid amounts from patient cost-sharing (copays and coinsurance).

[11] The pay-per-patient flag is observed at the claim level, indicating that the insurer pays the billing provider in this way. However, we do not observe the terms of the contract and can therefore not differentiate between different types of pay-per-patient contracts. Furthermore, we do not observe how individual physicians within a practice are remunerated (e.g., profit-sharing or salaried), leading us to assume that the contracting incentives are fully passed through. We therefore do not distinguish between the individual physician and their practice's incentives.

[12] Claim-level price information under pay-per-patient contracts is unreliable, as physicians operate under a budget rather than being paid per procedure. We therefore interpret prices under these contracts with caution.

[13] Our approach follows the logic of the TIN-matching method of Baker et al. (2016), which also aims to identify the business entity.

	Medicare Advantage		Commercial	
	Acquiring Insurer	Other Insurers	Acquiring Insurer	Other Insurers
Full Sample				
N	76,837	35,750	171,357	294,323
Age	69.64	68.97	43.15	43.37
	(0.03)	(0.04)	(0.03)	(0.03)
Share White	0.89	0.87	0.65	0.53
	(0.00)	(0.00)	(0.00)	(0.00)
Income	76,940	74,980	86,940	85,870
	(120)	(180)	(90)	(70)
Risk	0.84	1.04	0.34	0.35
	(0.00)	(0.01)	(0.00)	(0.00)
Claims	11.14	20.38	11.19	10.06
	(0.05)	(0.13)	(0.04)	(0.03)
Expenditure	3,282.33	5,897.08	4,612.45	4,668.28
	(49.40)	(74.26)	(49.02)	(37.83)
Acquired PCP Sample				
N	6,289	1119	10,351	14,939
Age	70.76	70.42	45.58	46.08
	(0.07)	(0.18)	(0.13)	(0.11)
Share White	0.96	0.97	0.70	0.58
	(0.00)	(0.01)	(0.01)	(0.01)
Income	90,060	98,800	94,330	94,970
	(420)	(1140)	(340)	(280)
Risk	0.70	0.94	0.36	0.37
	(0.01)	(0.03)	(0.00)	(0.00)
Claims	8.42	15.03	11.36	9.97
	(0.14)	(0.47)	(0.15)	(0.12)
Expenditure	2,470.46	4,125.98	3,742.85	4,249.30
	(93.29)	(345.31)	(143.04)	(169.91)

Table 1: Patient Descriptives By Segment, Insurer and Primary Care Practice in 2016

Notes: Reported values are averages, with standard errors in parentheses. We define beneficiaries of "Other Insurers" as the complement of beneficiaries of the "Acquiring Insurer". We proxy for income using zipcode-level average income from the Internal Revenue Service Statistics of Income (2016). Average risk scores are computed using CMS's 2019 software.

(3.1 physicians and 404 patients), practices at the upper end may care for over 50,000 patients and comprise more than 1,000 physicians.

We further define a subset of our sample as the *Acquired PCP Sample*, made up solely of patients and physicians assigned to the acquired practice. The acquired practice is in the 99^{th} percentile of the distribution, with about 80 assigned physicians and over 39,000 assigned patients each period. It is important to note that the acquisition does not imply exclusivity; that is, patients of other insurers may (and do) still visit the acquired practice.[14] Namely, the acquired practice continues to care for MA and Commercial beneficiaries of other insurers after the acquisition.

We present summary statistics for our sample in 2016 in Table 1. The acquiring insurer is the largest in the market, with market shares of 37% in the Commercial segment and of 68% in MA. The acquired practice is also the largest in our sample, and is the most common primary care provider for beneficiaries of the acquirer. Their patients are more likely to be White, have higher average income and have lower baseline risk scores.[15]

3 Conceptual Framework

Before proceeding with our empirical context, we theoretically illustrate how insurer-physician acquisition overcomes agency problems between insurers and primary care physicians using a simple conceptual framework. The purpose of this framework is to help guide and interpret our empirical analyses. Through the acquisition, PCPs are incentivized to increase their provision of diagnoses for MA beneficiaries, as this results in higher risk-adjustment payments. Furthermore, they modify their referral behavior to generate cost savings from specialist care.

Set Up. We focus on four types of agents in this setup: (1) the government, which sets risk adjustment payment rates for MA beneficiaries; (2) the insurer, which competes for beneficiaries and sets healthcare networks and prices; (3) the PCP, who makes joint decisions with patients on their healthcare; and (4) the specialists who provide care to patients referred by the PCP. The government, the insurer, and the specialists do not explicitly make any choices in our framework. Instead, we take the risk adjustment payment structure, network, and prices as given, and focus on the choices made by the PCP.[16,17]

The PCP's Problem. The PCP maximizes the joint surplus of their utility and patient health by choosing when to diagnose and where to refer patients. Patients belong to a market segment $\theta \in \{MA, C\}$, Medicare Advantage or Commercial, respectively, and have an underlying health condition, e.g., diabetes, $\delta^*_\theta \in \{0,1\}$.[18] The PCP's problem is divided into two stages. In the first stage, the PCP sees the patient, observes the probability that the patient is healthy p_θ, decides whether or not to carry out a diagnostic test, $e_\theta \in \{0,1\}$, and, lastly, decides whether to provide a

[14] In Appendix Figure E.2, we show the number of each insurer's patients assigned to the acquired practice by insurer and market segment. Appendix Table F.11 further evaluates whether there is steering into the acquired practice and shows an increase in the probability of the practice's patients to be MA beneficiaries of the acquiring insurer, as well as an increase in the probability of Commercial beneficiaries of other insurers.

[15] We proxy for income by using beneficiaries' zipcode's average income, computed by dividing a zipcode's total gross income by the number of returns in the zipcode. This measure is reported by the Internal Revenue Service Statistics of Income (2016).

[16] In a staged game version, the former objects would be the equilibrium outcomes of previous stages, e.g., of Nash-in-Nash bargaining for provider pricing and network inclusion, competition in the insurance market for premiums, etc. However, this is not central to our argument, so we abstract away from it for the sake of simplicity.

[17] Our set up is similar to that of **?**, but we focus on physician incentives and behavior while abstracting away from patient sorting into MA and Traditional Medicare.

[18] We consider one diagnosis for the sake of simplicity, but the intuition carries for N conditions.

diagnosis, $d_\theta \in \{0,1\}$. Then, in a second stage, if a patient is diagnosed, the PCP may refer them to a specialist for care. The PCP chooses referral intensity $r_\theta \in \mathbb{R}^+$ to maximize their joint utility with their patients $v_\theta^{\delta^*}(r_\theta)$.

In the first stage, the PCP can avoid uncertainty in patient health p_θ by ordering a diagnostic test e_θ. If the PCP chooses not to carry out any tests, i.e., $e_\theta = 0$, they decide whether or not to diagnose the patient as a function of the patient's probability of being healthy p_θ. If, instead, the PCP carries out the diagnostic test, i.e., $e_\theta = 1$, they perfectly observe their patient's health status and diagnose only unhealthy patients. We then write the PCP's payoffs under each combination of diagnostic effort and diagnostic decisions as follows:

$$u^{PCP}(e_\theta, d_\theta, r_\theta) = \begin{cases} (1-p_\theta)\alpha_1 & \text{if } e_\theta = 0 \ \& \ d_\theta = 0 \\ p_\theta\alpha_0 + \beta_\theta + p_\theta v_\theta^0(r_\theta) + (1-p_\theta)v_\theta^1(r_\theta) & \text{if } e_\theta = 0 \ \& \ d_\theta = 1 \\ (1-p_\theta)(\beta_\theta + v_\theta^1(r_\theta)) - (\omega - \gamma_\theta) & \text{if } e_\theta = 1 \end{cases} \tag{1}$$

Here, the PCP incurs a cost α_1<0 from failing to diagnose an unhealthy patient, and a cost α_0 from a false positive diagnosis.[19,20] They may further be paid at a rate $\beta_\theta \geq 0$ for providing a diagnosis, which allows for there to be no payment from diagnosing ($\beta = 0$), for instance, under a fee-for-service payment arrangement, as well as positive payment, e.g., under pay-per-patient contracts. Lastly, the PCP pays an effort cost ω to conduct diagnostic tests, and may be reimbursed γ_θ by the insurer.

In the second stage, we model the joint utility of the PCP and their patients for referrals in the function $v_\theta^{\delta^*}(r_\theta)$, assumed to be increasing and concave. This utility is designed to represent both the value of referrals to patients and the effort cost for the PCP to provide referrals. Patients' value of referrals is allowed to vary based on their true underlying health status; that is, patients may value referrals more when they are unhealthy and specialist care is more beneficial to them than when they are healthy. The PCP's cost of providing a referral can be thought of as the cost of searching for a good specialist match within their patients' insurance network. Lastly, note that we express the PCP's problem in money-metric units.

Other Agents' Revenues and Costs. Although other agents do not make choices in our framework, we model their payoffs below. We start by modeling the insurer's profits. Their revenues consist of commercial premiums and MA risk-adjustment payments, and their costs arise from care provision. We normalize demand in each segment to 1 to simplify notation. We further omit premiums in MA as they are approximately zero in our setting, and define premiums in the Commercial segment as p^C. Then, we define risk adjustment payments to the insurer as $p^R d_{MA}$, where p^R is the rate set by the government. Lastly, insurers incur costs $\beta_\theta d_\theta$ and $\gamma_\theta e_\theta$ to compensate the PCP for their diagnostic provision and effort, respectively, as well as specialist costs $s_\theta(r_\theta)d_\theta$ to

[19]Note that, as we only observe realized diagnoses in our data, we cannot measure these quantities.
[20]We follow Chan et al. (2022), setting up the PCP's cost of misdiagnosis as a problem of statistical classification.

compensate specialists for their provision of care. The insurer's profits are then:

$$\Pi = p^C + p^R d_{MA} - \sum_{\theta} \left(\beta_\theta d_\theta + \gamma_\theta e_\theta + s_\theta(r_\theta) d_\theta \right),$$

where we note that insurer's profits at this stage do not depend on insurer choices or demand in the insurance market, but solely on the PCP's care decisions.

Next, we model specialists, who incur costs and receive payments from the insurer for their provision of care. Their payoffs are: $(s_\theta(r_\theta) - c_\theta(r_\theta))d_\theta$, where we note that specialists only receive referrals and therefore only provide care if patients have been diagnosed. Lastly, we define the government's excess burden of raising public funds as κ. This yields government expenditures: $(1+\kappa)p^R d_{MA}$.

Welfare. We combine the previous expressions to express a welfare function as follows:

$$W_\theta(e_\theta, d_\theta, r_\theta) = \begin{cases} \overbrace{(1-p_\theta)\alpha_1}^{\text{Misdiag. cost}} & \text{if } e_\theta = 0 \ \& \ d_\theta = 0 \\ \overbrace{p_\theta \alpha_0}^{\text{Misdiag. cost}} + \overbrace{\mathbb{E}_{p_\theta}(v_\theta(r_\theta)) - c_\theta(r_\theta)}^{\text{Net referral value}} - \overbrace{\kappa p^R \mathbb{1}\{\theta = MA\}}^{\text{Public funds burden}} & \text{if } e_\theta = 0 \ \& \ d_\theta = 1 \\ (1-p_\theta) \left(\underbrace{v_\theta^1(r_\theta) - c_\theta(r_\theta)}_{\text{Net referral value}} - \underbrace{\kappa p^R \mathbb{1}\{\theta = MA\}}_{\text{Public funds burden}} \right) - \underbrace{\omega}_{\text{Effort Cost}} & \text{if } e_\theta = 1 \end{cases} \tag{2}$$

Although we do not estimate this function, it is useful to model welfare to understand the different impacts that the acquisition may have. On the diagnostic side, welfare depends on the cost of misdiagnosis, and diagnostic effort, as well as the excess burden of public funds incurred to pay out the risk adjustment to the insurer. On the referral side, welfare depends on the net value of referrals, that is the utility of referrals $v_\theta(r_\theta)$ net of the referral cost $c_\theta(r_\theta)$.

Baseline Equilibrium. Because referrals require a diagnosis, we solve this sequentially. First, we solve for the optimal referral intensity given effort and diagnostic choices. Then, given referrals, we solve for diagnostic decisions. We start by solving for referrals under $e_\theta = 0$ and $d_\theta = 1$ as follows:

$$\max_{r_\theta} \ v_\theta^0(r_\theta) p_\theta + v_\theta^1(r_\theta)(1-p_\theta)$$

We then compute the first-order condition for referrals as follows, where r_θ^* represents the optimal referral choice:

$$\frac{v_\theta^{1\prime}(r_\theta^*(e_\theta = 0))}{v_\theta^{0\prime}(r_\theta^*(e_\theta = 0))} = -\frac{p_\theta}{1-p_\theta}$$

This captures how, when the PCP faces uncertainty over their patient's health status, they optimally choose their referrals so that the ratio of marginal utility from the extra referral between states equals the inverse ratio of the probabilities of the patient being in each health state.

Lastly, we solve for the optimal referral intensity under $e_\theta = 1$. In this case, the PCP perfectly observes the patient's health status and may only refer a diagnosed patient. Their payoff from referral is then $v_\theta^1(r_\theta)$, and the optimal referrals satisfy $v_\theta^{1\prime}(r_\theta^*(e_\theta = 1)) = 0$. Note that the optimal choice of referrals is higher when the PCP exerts diagnostic effort, that is, $r_\theta^*(e_\theta = 1) > r_\theta^*(e_\theta = 0)$.

Next, we solve for the PCP's diagnostic choices, taking their referral choices as given by the above first-order conditions. The PCP chooses whether to exert effort and to diagnose their patient, with observed probability of being healthy p_θ, to maximize their utility in Equation 1. Their optimal decisions are characterized by thresholds of observed probability of their patient's health. We further assume that the PCP's utility under each effort and diagnostic choice are monotonic in their patients observed health signal. We derive the expression for the decision rule in Appendix B, and summarize it as follows:

$$\{e_\theta^*, d_\theta^*\} = \begin{cases} \{0,1\} & \text{if } p_\theta < \underline{p}_\theta^* \\ \{1,\cdot\} & \text{if } \underline{p}_\theta^* \leq p_\theta \leq \bar{p}_\theta^* \\ \{0,0\} & \text{if } p_\theta > \bar{p}_\theta^* \end{cases} \tag{3}$$

Then, for p_θ large enough, the PCP would optimally choose to not exert any effort and not diagnose the patient, whereas if it is low enough they would choose to not exert any effort and diagnose their patient regardless. Therefore, the PCP only exerts effort if they are uncertain enough about their patient's true underlying health status.

The Acquisition and Incentive Alignment. The insurer can align the PCP's incentives with their own through acquisition or through contracting. We focus next on how acquisition aligns incentives, while noting that our predictions match those under the adoption of a pay-per-patient contract. We discuss the latter in Appendix B.

We model the acquired PCP as choosing care provision to maximize a joint utility function. In this way, the PCP incorporates both the risk-adjustment payments to the insurer when a patient is diagnosed, as well as the costs of care provision by both specialists and themselves. We further include an incentive alignment friction $\lambda \in [0, 1]$ to account for imperfect incentive alignment, where $\lambda = 1$ represents perfect incentive alignment and $\lambda = 0$ if there is no incentive alignment.

$$u^{ACQ\ PCP}(\cdot|\lambda) = \begin{cases} u^{PCP}(0,0,\cdot) & \text{if } e_\theta = 0 \ \& \ d_\theta = 0 \\ u^{PCP}(0,1,r_\theta) + \lambda \left(p^R \mathbb{1}\{\theta = MA\} - \beta_\theta - s_\theta(r_\theta)\right) & \text{if } e_\theta = 0 \ \& \ d_\theta = 1 \\ u^{PCP}(1,\cdot,r_\theta) + (1-p_\theta)\lambda(p^R \mathbb{1}\{\theta = MA\} - \beta_\theta - s_\theta(r_\theta)) - \lambda\gamma_\theta & \text{if } e_\theta = 1 \end{cases}$$

Our model generates two predictions for how referrals and diagnoses change after the acquisition. We start by discussing the first prediction, on the impact of the acquisition on referrals:

Prediction 1. *The acquisition (almost always) leads to a decrease in referrals.*

Intuitively, for patients who are eligible to be referred before and after the acquisition, the PCP decreases the intensity of referrals as they internalize the cost of specialist care. Furthermore, this decrease in referrals is larger as incentives become more aligned and the PCP internalizes a larger share of the specialist cost. Lastly, although some patients become eligible to be referred after the acquisition, insofar as the marginal cost of referrals is larger than the marginal benefit of referrals among healthy patients, referrals are still decreasing. We provide a detailed derivation of these results in Appendix B.

Next, we discuss the second prediction, on the impact of the acquisition on expected diagnoses:

Prediction 2. *The acquisition generates an increase in expected diagnoses. We define this increase as differential coding.*

Intuitively, this is the result of the acquisition increasing the value of providing diagnoses for the PCP, who internalizes both the risk-adjustment payments to the insurer and the cost of payments to the PCP. To see this, we solve for the PCP's diagnostic choices as before, which yields a new set of cutoffs, $\{\underline{p}_\theta^{\mathcal{I}}, \bar{p}_\theta^{\mathcal{I}}\}$. We derive and present these analytical results in Appendix B, and discuss their intuition below.

First, we show that the lower bound is increasing in the degree of incentive alignment. Intuitively, as the cost of effort and the value of diagnoses both increase, the PCP is more likely to directly diagnose their patients and not exert effort among patients who are more likely to be unhealthy. Second, we examine the impact on the upper bound, where exerting effort is needed to diagnose the marginal patient. Then, in so far as the cost of effort is outweighed by the benefits of diagnosing the marginal patient, the upper bound is increasing in incentive alignment. Lastly, note that, whereas the impact on the lower bound decreases effort, the impact on the upper bound increases effort, yielding ambiguous implications for the total impact on diagnostic effort.

Finally, we evaluate how the acquisition impacts the ex-ante probability of diagnosing a patient, where we assume the observed probability of a patient's health status is drawn from a uniform distribution. We start by defining the probability of providing a diagnosis as follows:

$$\mathbb{E}_{p_\theta}(d_\theta) = \underline{p}_\theta + (\bar{p}_\theta - \underline{p}_\theta)\left(1 - \frac{\bar{p}_\theta + \underline{p}_\theta}{2}\right)$$

The impact of incentive alignment on the probability of diagnosis is then determined by the following equation:

$$\frac{\partial \mathbb{E}_{p_\theta}(d_\theta)}{\partial \lambda} = \frac{\partial \bar{p}_\theta}{\partial \lambda}(1 - \bar{p}_\theta) + \frac{\partial \underline{p}_\theta}{\partial \lambda}\underline{p}_\theta > 0,$$

where we note that, regardless of the total impact on effort, an increase in incentive alignment generates an increase in the probability of diagnosis.

Welfare Impacts. Given these changes in the physician's decisions, we discuss potential implications for welfare. We leverage our expression for welfare in Equation 2 to outline the different impacts of the acquisition by computing the change in welfare before and after the acquisition. Below, we present a stylized version of our expression, and present its component parts in detail in Appendix Equation B.14.

$$\Delta W_\theta = \Delta\text{Diagnostic precision} - \Delta\text{Excess burden of public funds} - \Delta\text{Effort cost} + \Delta\text{Referral utility} \tag{4}$$

We begin by examining the welfare implications of increased diagnoses in the first three terms. The first term captures the change in welfare from changes in the precision of diagnostic behavior, although we note that we cannot empirically speak to this impact as we don't observe patients' true underlying health status. On the one hand, increasing the lower bound to diagnose more patients without testing, increases diagnoses among healthy patients, generating a welfare loss. On the other hand, increased testing among patients who present as healthier increases diagnoses among sick patients who appeared healthy. The overall impact of this depends on how we trade off between over- and under-diagnosing, that is, the relative magnitudes of α_0 and α_1, as well as the relative sizes of each of these populations.[21] The second term captures the loss in welfare from the excess burden of government expenditures as the PCP provides more diagnoses. Lastly, the third term captures the change in welfare from the change in diagnostic effort. Whether this increases or decreases welfare depends on the total effect on effort, which is ambiguous in our model.

Then, we examine the impact on welfare from changes in referral behavior in the fourth term. In this term, we trade off the extensive margin change of referring more individuals as they are

[21] This depends on the magnitude of the increase in the cutoff at each end, and, under different distributions, this also depends on the mass of people near each cutoff.

diagnosed and become eligible, and the intensive margin of reducing referrals after integration. Insofar as the reduction in referrals generates larger cost savings relative to the decreased patient utility, then the intensive margin reduction in referrals is beneficial for welfare. Furthermore, note that the scope for this to be beneficial depends on the variation in specialist cost. Lastly, the increase in the extensive margin may be beneficial or detrimental for welfare for different groups of patients. Among healthy patients, this decreases welfare, whereas it improves it for unhealthy patients.

4 Acquisitions and Risk Adjustment Payments

In this section, we empirically evaluate whether the acquisition delivers increased diagnoses in Medicare Advantage, as illustrated by our framework. We find that the acquired physician practice increased their diagnostic provision, delivering average increased payments of $1508 per patient per year in MA. We start by examining some descriptive patterns. Then, we outline the empirical strategy we leverage to estimate the impact of the acquisition and present our estimates. Lastly, we discuss the interpretation and normative implications of our results.

4.1 Descriptive Findings

To evaluate the impact of the acquisition on risk adjustment payments to the insurer, we start by exploring patient-level patterns over time. We inspect two outcomes: first, whether patients are covered under a pay-per-patient contract at their primary care practice; and second, the patient's risk score as computed with the CMS software. As both of these outcomes are determined on a yearly basis, we define an observation as a patient-year.

Our model suggests that acquisition and pay-per-patient contract adoption may both align incentives between insurers and physicians to deliver increased diagnoses. We then start by inspecting the prevalence of these contracts at both non-acquired practices and the acquired practice, where contract adoption can complement the acquisition. It is further relevant to understand contract adoption in the market, as it may help determine the appropriate control group for evaluating the impacts of the acquisition.

Then, examining the prevalence of pay-per-patient contracts in the market reveals widespread adoption among Medicare Advantage patients of the acquiring insurer. We define a patient as being covered under one of these contracts in a given year if they have at least one claim at their assigned PCP that was paid under a pay-per-patient contract. We find that, co-timed with the acquisition, the acquiring insurer adopts pay-per-patient contracts with primary care practices covering approximately 50% of their MA beneficiaries (Figure 1, top-left panel). In contrast, there is no adoption in the Commercial segment (Figure 1, top-right panel).[22] Lastly, these contracts

[22] A potential concern is that, after acquisition, claims are flagged as pay-per-patient for internal firm purposes rather than because PCPs are actually being remunerated through one of these contracts. However, the fact that this is not true for Commercial beneficiaries of the acquirer at the acquired practice alleviates these concerns. That is, if this were just an artifact of changes in internal claims reporting, we would see the same change across the board.

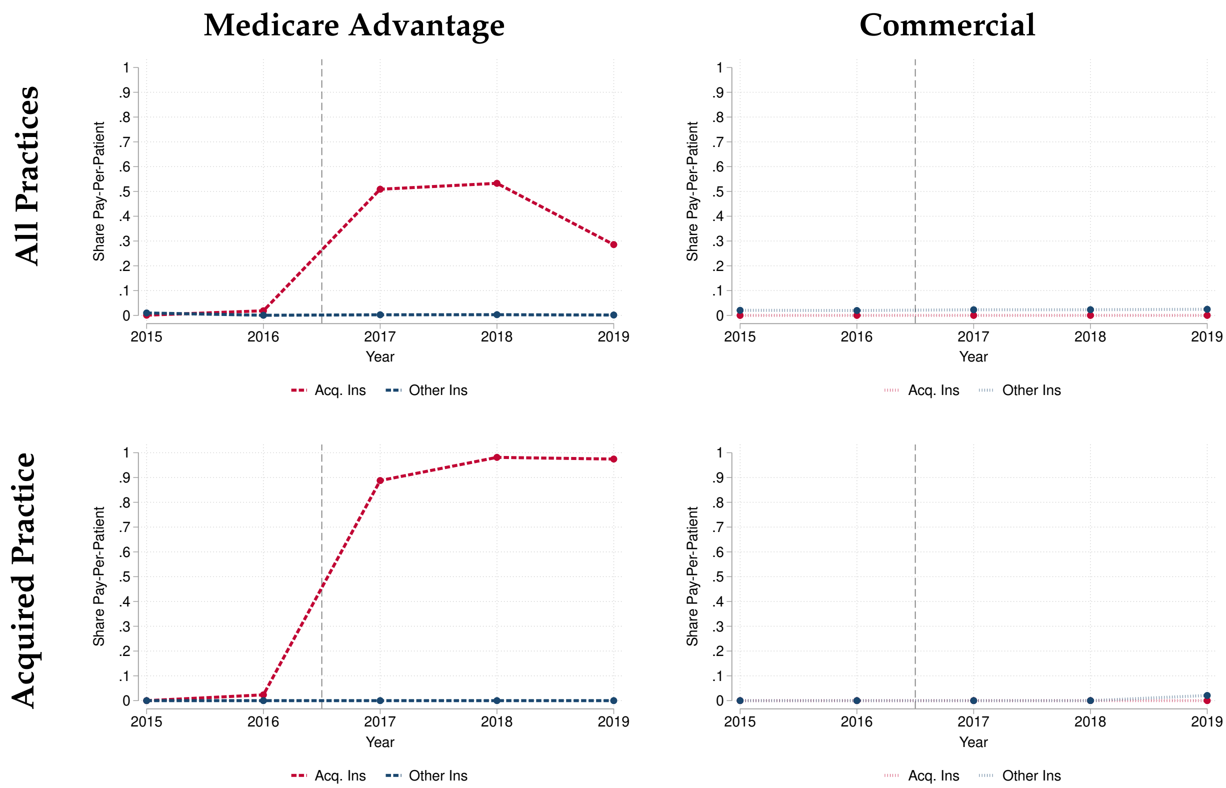

Figure 1: The Acquiring Insurer Adopts Pay-Per-Patient Contracts in MA

Notes: The figures illustrate the acquirer's adoption of pay-per-patient contracts, which occurs exclusively in the Medicare Advantage (MA) market. While the insurer implements capitation for 40-50% of its MA beneficiaries across its entire network (top-left), the policy is applied to nearly 100% of its MA beneficiaries at the acquired practice post-acquisition (bottom-left). The vertical line in each panel indicates the 2017 acquisition date. We define a patient as being covered under a pay-per-patient contract if they have at least one claim at their assigned PCP in a given year that was paid under this arrangement. We then compute the share of patients under one of these contracts by dividing the number of patients under a pay-per-patient contract at their PCP by the total number of patients assigned to a PCP from a given insurer.

are adopted for virtually all of the insurer's MA beneficiaries at the acquired practice (Figure 1, bottom-left panel).

Then, we examine patient risk scores over time and find an upward trend in the risk scores of MA beneficiaries of the acquiring insurer over our sample period. We compute patient risk scores by inputting patient diagnoses and demographic information into the 2019 CMS-HCC software (as described in Section 2.2) and plot patients' risk by insurer and market segment in Figure 2. This figure illustrates how the risk scores of MA beneficiaries of the acquiring insurer increase over the course of our sample, both at the acquired practice and elsewhere, whereas there is no discernible change in the risk scores of Commercial beneficiaries.

These patterns are consistent with the predictions arising from our framework. In MA, insurers can align incentives to increase risk scores and, consequently, risk adjustment payments

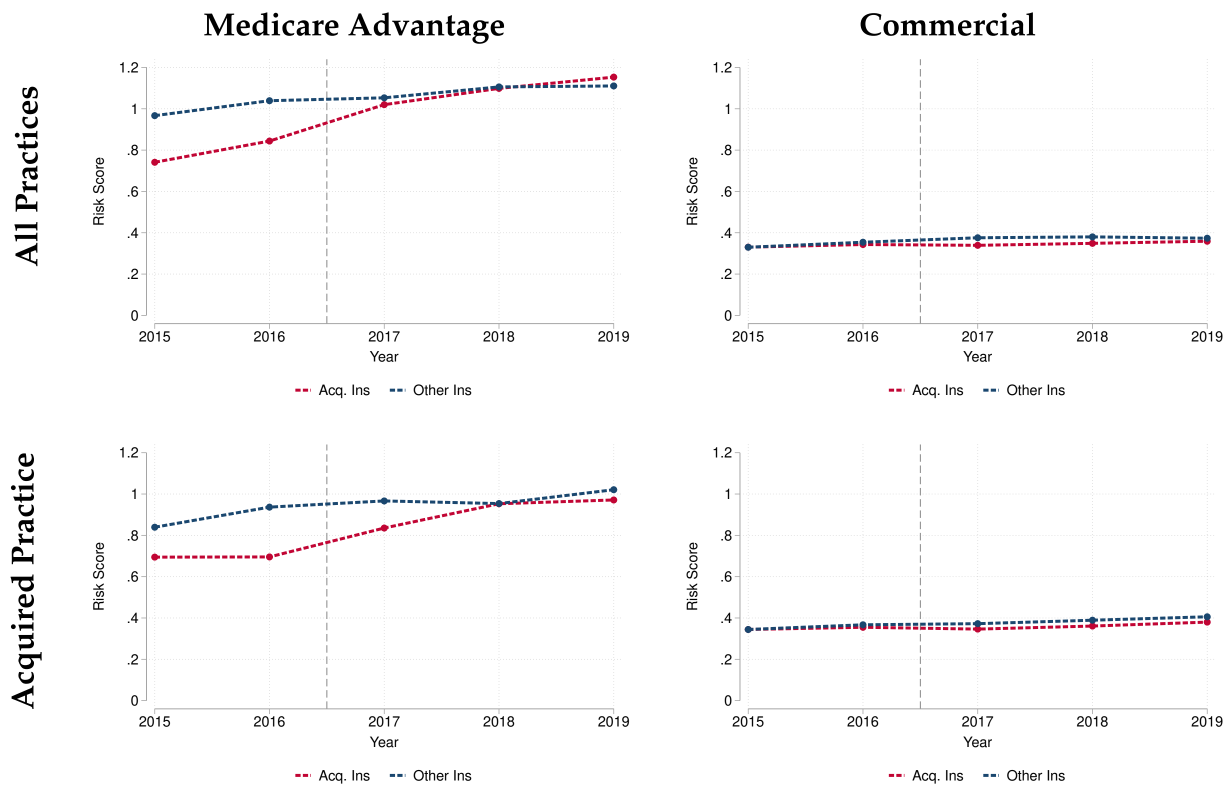

Figure 2: Average Risk Score Over Time Across Groups

Notes: The figures illustrate how the risk scores of MA beneficiaries of the acquiring insurer increase over the course of our sample, both at the acquired practice and elsewhere, whereas there is no discernible change for the risk scores of Commercial beneficiaries. Risk scores are computed using the 2019 CMS-HCC software. The vertical line in each panel indicates the 2017 acquisition date.

through both acquisition and adoption of pay-per-patient contracts. In Commercial, as patient risk scores do not impact insurer revenues, there are no incentives to increase reported risk. We therefore restrict attention to the MA segment for the rest of the analysis of the impact on risk adjustment payments.

4.2 Empirical Strategy

Our objective is to measure how the acquisition impacts physicians' diagnostic behavior. That is, whether the acquired practice changes their diagnostic behavior to deliver increased risk adjustment payments to the acquiring insurer. However, the descriptive patterns above highlight the key challenges to identifying this impact. First, the endogeneity of acquisition creates concerns over the selection of the acquired practice into the acquisition. Second, comparisons to other practices are contaminated by their adoption of pay-per-patient contracts with the acquiring insurer, as shown in Figure 1.

We summarize our strategy to overcome these challenges as restricting attention to our *Acquired PCP Sample* and comparing MA patients of the acquiring insurer at the acquired practice to MA patients of other insurers at the acquired practice. Note that, because the acquisition is co-timed with the adoption of a pay-per-patient contract, our strategy captures the bundled treatment effect on the acquiring insurer's patients' risk scores. We further leverage the panel structure of our data to include individual fixed effects and measure within-patient changes in risk score. The identification assumption is that, absent the acquisition, the acquiring insurer's patients' risk scores would have evolved in the same manner as those of other insurers' patients at the acquiring PCP.

Our full setup is described in the causal graph in Equation 5. Let i denote a patient and j a PCP. The vector $\mathbf{A}$ indicates the acquisition status, where A_i is an indicator for whether patient i is a beneficiary of the acquiring insurer, and A_j is an indicator for whether the PCP was acquired. Furthermore, define $PC(\mathbf{A})$ as an indicator for pay-per-patient contract adoption as a function of acquisition, and $R(\mathbf{A}, PC)$ as the potential risk score for a patient given acquisition and pay-per-patient contract adoption by their PCP. Then, our treatment group and control group, as defined above, have observed outcomes $R(1,1,1)$ and $R(0,1,0)$, respectively.

$$\mathbf{A} = (A_i, A_j) \to PC(A_i, A_j) \to R(A_i, A_j, PC) \tag{5}$$

Furthermore, because our objective is to measure changes in the behavior of the acquired practice, we define our main outcome as the patients' risk score resulting solely from diagnoses provided by the acquired practice, denoted by $R^{ACQ\ PCP}$. In this way, we isolate the change in risk score arising solely from diagnostic decisions by the acquired practice. In contrast, R^{ALL}, the risk score computed using all diagnoses for a patient-year, may conflate changes in the acquired practice's behavior with changes in other practices' behaviors.

Finally, we estimate the impact of the acquisition on reported risk scores by conducting an event study design with individual and time fixed effects following Callaway and Sant'Anna (2021).[23] Our estimates capture how the acquisition leads the acquired PCP to treat patients of the acquirer differently from those of other insurers. Note that, in the presence of within-practice spillovers onto other patients, our estimates are a conservative estimate of the impact. Lastly, because pay-per-patient contract adoption is co-timed with the acquisition, we can only estimate the reduced form impact of the acquisition. We discuss interpretations of how the acquisition and pay-per-patient contracts interact in Section 4.4.1.

4.3 Results

4.3.1 Differential Coding from the Acquisition

We find that, after the acquisition, the acquired practice provides increased diagnoses for the MA beneficiaries of the acquiring insurer. On average, this change in behavior results in increased risk

[23] Because individuals rarely switch insurers in our sample, the individual fixed effects also capture insurer effects. This is not uncommon more broadly, indeed, 78% of MA beneficiaries stay in the same plan as shown by Jacobson et al. (2016).

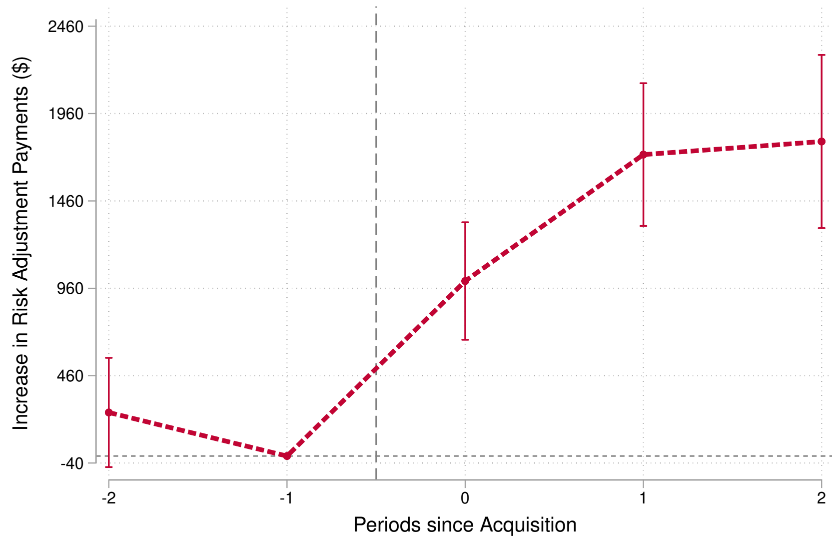

Figure 3: Differential Coding at the Acquired PCP

Notes: The figure presents the event study estimates for the change in risk adjustment payments to the acquiring insurer per patient per year. We find an increase in payments of $998 to $1805. The vertical line indicates the 2017 acquisition date. An observation is a patient-year. We restrict the sample to our Acquired PCP sample, that is, we keep only patients of the acquired practice. The treated group are MA beneficiaries of the acquiring insurer at the acquired practice. The control group are MA beneficiaries of other insurers at the acquired practice. We further restrict attention to diagnoses provided by the acquired practice. Risk scores are converted into dollar payments using the constant for Denver in 2017. Our estimates are computed using the event study design with individual and time fixed effects from Callaway and Sant'Anna (2021).

adjustment payments of $1508 per patient per year. We present these event study estimates in Table 2, column (1), and in Figure 3, where we convert patients' risk scores to dollar payments.[24]

By restricting the risk score to diagnoses provided by the acquired PCP, our main estimates capture the impact of treatment directly on the acquired practice's diagnostic behavior. However, this impact could either under- or overestimate the full impact of the acquisition on the insurer's risk adjustment payments. On the one hand, if the acquisition led to a "business stealing" effect, whereby the acquired PCP replaced other physicians in providing diagnoses but did not provide any diagnoses that a different physician would not have provided otherwise, payments to the insurer would not change, and we would be overestimating the impact. On the other hand, because providing diagnoses can be costly, the acquired practice may avoid duplicating diagnoses, in which case our measured impact could be smaller than the full impact on payments to the insurer.

We evaluate whether our main estimates are a lower or upper bound on the full impact by conducting the same event study design while utilizing the risk score from all diagnoses (R^{ALL})

[24]We do this conversion using the payment rate for Denver in 2017, of approximately $800 per month (Source: CMS 2017 Medicare Advantage and Prescription Drug rate information).

	(1) $R^{ACQ\ PCP}$	(2) R^{ALL}	(3) R^{ALL}
τ_{-2}	0.026	0.111	-0.011
	(0.017)	(0.026)	(0.026)
τ_{-1}	0	0	0
τ_0	0.104	0.125	-0.30
	(0.018)	(0.032)	(0.030)
τ_1	0.180	0.233	0.013
	(0.022)	(0.039)	(0.063)
τ_2	0.188	0.224	-0.050
	(0.026)	(0.041)	(0.069)
Individual FE	X	X	X
Year FE	X	X	X
Baseline Mean	0.514	0.696	0.937
Observations	28,204	28,204	136,813
PCP Sample	ACQ PCP	ACQ PCP	All
Insurer Sample	All	All	All but ACQ INS

Table 2: Integration and Differential Coding

Notes: The table presents the event study estimates for the change in patient risk scores across three specifications. An observation is a patient-year. In the first two specifications, we restrict the sample to our Acquired PCP sample, that is, we keep only patients of the acquired practice. The treated group are MA beneficiaries of the acquiring insurer at the acquired practice. The control group are MA beneficiaries of other insurers at the acquired practice. Columns (1) and (2) differ in the diagnoses used for the risk score construction. In column (1), we restrict to diagnoses provided by the acquired practice. In column (2), we use all diagnoses. Column (3) considers all PCP practices but restricts to beneficiaries of other insurers. This sample requires that we use all diagnoses to compute the risk score. Our estimates are computed using the event study design with individual and time fixed effects from Callaway and Sant'Anna (2021). Control is not yet treated. Baseline mean is the mean of the treated group at $t = -1$.

as our outcome. We find, in Table 2, column (2), that the impact of the acquisition on patients' full risk is slightly larger, if not significantly different, than that of our main specification. This rules out that our preferred specification is overestimating the full impact, and, if anything, suggests that the acquired PCP avoids re-duplicating diagnoses provided by other providers. We further evaluate the magnitude of our estimates by comparing the magnitude of our estimates to existing estimates. Geruso and Layton (2020) find differential coding ranging from .06 to .16 risk points, where provider owned plans are on the higher end. Furthermore, MedPac estimate a .08 increase in 2018, increasing to a projected .20 increase in 2025. Then, Our estimates are on the higher end of existing estimates of differential coding, ranging from .1 to .19 risk points.

Furthermore, our main specification (Table 2, column (1)) captures how the acquired practice differentially treats patients of the acquiring insurer as compared to those of other insurers. However, the acquisition may have spillover effects on patients of other insurers. For instance, changes in administrative procedures could have practice-wide impacts, potentially increasing diagnoses across patient groups. Then, our estimates would underestimate the impact of the acquisition. We test for practice-wide effects by comparing MA beneficiaries of other insurers at the acquired practice, that is the control group in our main specification, to MA beneficiaries of other insurers who are not patients of the acquired practice using our *Full Sample*.[25] We find no significant impact of the acquisition on the risk scores of MA beneficiaries of other insurers at the acquired practice (Table 2, column (3)), alleviating these concerns.

A further concern arises if the acquiring insurer's beneficiaries face systematically different trends in risk scores from other beneficiaries at the acquiring PCP. While it is true that patients of the acquiring insurer at the acquired PCP have lower risk scores than those of other insurers in the pre-period, our treatment and control groups do appear demographically comparable (Table 1, Acquired PCP Sample). Furthermore, the inclusion of individual fixed effects should partially alleviate these concerns as we compute individual-level changes in risk. This specification also addresses issues with selection into the acquired practice after the acquisition, which may arise if, for instance, patients with more room for increased diagnoses are steered into the practice after the acquisition. We present specifications without individual fixed effects in Appendix Table F.2, which slightly attenuate our estimates. Lastly, as we only have two years of data before the acquisition and risk scores are an annual object, we only have two pre-treatment periods.

4.3.2 Differential Coding from the Pay-per-Patient Contract Adoption

Next, we examine how defining treatment as the adoption of a pay-per-patient contract changes our estimates. We interpret our main specification as capturing the reduced-form impact of the acquisition, which, conceptually, can operate through the adoption of pay-per-patient contracts. Then, using the adoption of pay-per-patient contracts captures their direct impact, i.e., the second arrow in Equation 5.

First, we focus on the acquired practice, where we define our treated group as those MA patients of the acquired practice for which the acquired practice has adopted a pay-per-patient contract with the acquiring insurer, and our control group as all MA patients who have not adopted a pay-per-patient contract.[26] Note that, because virtually all MA beneficiaries of the acquiring insurer and none of the beneficiaries of other insurers at the acquired practice adopt one of these contracts, we expect the impacts to be qualitatively similar. Nevertheless, our estimates may differ for two reasons: first, adoption is slightly staggered in our sample, and second, adoption is a slightly noisier measure of treatment than the acquisition. Indeed, in Table 3, columns (1) and (2), we find that defining treatment as pay-per-patient adoption yields qualitatively similar estimates

[25] Under our potential outcomes framework, our treated outcome is then $R(0,1,0)$ and the control outcome is $R(0,0,0)$. Furthermore, note that, to conduct this event study, we must use all diagnosis risk, R^{ALL}, as our outcome.

[26] Under our potential outcomes framework, our treated outcome is then $R(1,1,1)$ and control outcomes are $R(1,1,0)$ and $R(0,1,0)$, for MA beneficiaries of the acquirer who are not yet under a pay-per-patient contract and MA beneficiaries of other insurers, respectively.

	(1) $R^{ACQ\ PCP}$	(2) R^{ALL}	(3) R^{ALL}	(4) R^{ALL}
τ_{-4}	0.275	0.438	-0.109	-0.101
	(0.069)	(0.088)	(0.037)	(0.039)
τ_{-3}	0.037	0.123	0.054	0.078
	(0.024)	(0.047)	(0.009)	(0.011)
τ_{-2}	0.014	0.065	-0.038	-0.038
	(0.009)	(0.015)	(0.005)	(0.005)
τ_{-1}	0	0	0	0
τ_{0}	0.107	0.101	0.150	0.134
	(0.011)	(0.019)	(0.005)	(0.005)
τ_{1}	0.094	0.128	0.165	0.153
	(0.022)	(0.037)	(0.006)	(0.007)
τ_{2}	0.080	0.142	0.182	0.155
	(0.037)	(0.066)	(0.009)	(0.011)
τ_{3}			0.315	0.325
			(0.048)	(0.051)
Individual FE	X	X	X	X
Year FE	X	X	X	X
Baseline Mean	0.510	0.665	0.858	0.858
Observations	31,616	31,616	561,933	372,372
PCP Sample	ACQ PCP	ACQ PCP	All but ACQ	All but ACQ
Insurer Sample	All	All	All	Only ACQ

Table 3: Differential Coding Across Samples

Notes: The table presents the event study estimates for the change in patient risk scores across four specifications. An observation is a patient-year. In the first two specifications, we restrict the sample to our Acquired PCP sample, that is, we keep only patients of the acquired practice. The treated group are MA beneficiaries of the acquiring insurer under a pay-per-patient contract at the acquired practice. The control group are MA beneficiaries of other insurers at the acquired practice as well as those MA beneficiaries of the acquiring insurer not yet under a pay-per-patient contract. Columns (1) and (2) differ in the diagnoses used for the risk score construction. In column (1), we restrict to diagnoses provided by the acquired practice. In column (2), we use all diagnoses. Columns (3) and (4) considers all PCP practices except for the acquired practice. In both columns, the treated group is made up of MA beneficiaries of the acquiring insurer whose assigned PCP adopted a pay-per-patient contract. The control group differs across columns. In column (4) we restrict to those MA beneficiaries of the acquiring insurer who have not yet or never adopt a pay-per-patient contract. In column (3) we include both this population and MA beneficiaries of other insurers. These sample requires that we use all diagnoses to compute the risk score. Our estimates are computed using the event study design with individual and time fixed effects from Callaway and Sant'Anna (2021). Control is not yet treated. Baseline mean is the mean of the treated group at $t = -1$.

as the acquisition treatment.[27]

Lastly, we shift our focus to the impact of pay-per-patient contracts in the absence of acqui-

[27] Slight differences in estimates and sample size arise from staggered adoption and the necessary balance for individual fixed effects when using the Callaway and Sant'Anna (2021) event study design.

sition and find similar impacts on risk-adjustment payments to the insurer. In Table 3, column (3), we present the event study estimates of the impact of pay-per-patient contract adoption on the risk scores of MA beneficiaries of the acquiring insurer. We exclude the acquired practice, use all diagnosis risk (R^{ALL}) as our outcome, and define our treated group as those MA beneficiaries of the acquirer covered under a pay-per-patient contract at their PCP, and our control group as MA beneficiaries of the acquirer who are not under a pay-per-patient contract as well as MA beneficiaries of all other insurers.[28]

We find that, on average, pay-per-patient contract adoption yielded a 0.203 increase in total risk in the post-periods, or an increase in risk adjustment payments to the insurer of approximately $1,949 per patient per year. Furthermore, restricting to within insurer comparisons, that is, dropping beneficiaries of other insurers from our control group, yields an average increase in payments of $1,840 in the post periods, not significantly different from the full estimate (Table 3, column (4)). However, there is substantial heterogeneity in the former across cohorts. In Table 4, we present estimates of the impact of pay-per-patient contract adoption by year of contract adoption, which reveal that early adopters of these contracts exhibit significantly larger impacts than those who adopt later.

4.3.3 Margins of Differential Coding

There are different ways in which physicians may increase their patients' reported risk. The mechanism used to increase patients' risk scores can matter for how we interpret differential coding and how well we think it tracks patients' true underlying risk. Because risk scores are a weighted average of patients' diagnoses, these mechanisms can operate on the extensive margin, by providing patients with any diagnosis, or on the intensive margin, by providing patients with more diagnoses or diagnoses with higher weights, conditional on having at least one diagnosis. We define these margins as follows, where D_{it} represents the number of diagnoses received by patient i in period t and R_{it} is their risk score:

$$\text{Any Diagnosis} = \mathbb{1}\left\{D_{it}^{ACQ\ PCP} > 0\right\} \tag{6}$$

$$\text{Quantity} = D_{it}^{ACQ\ PCP}\mathbb{1}\left\{D_{it}^{ACQ\ PCP} > 0\right\} \tag{7}$$

$$\text{Severity} = \frac{R_{it}^{ACQ\ PCP}}{D_{it}^{ACQ\ PCP}}\mathbb{1}\left\{D_{it}^{ACQ\ PCP} > 0\right\} \tag{8}$$

On the extensive margin, insurers and PCPs are anecdotally more likely to reach out to patients for a yearly check-up, or even send home health aides to visit patients and provide diag-

[28]Under our potential outcomes framework, our treated outcome is then $R(1,0,1)$ and control outcomes are $R(1,0,0)$ and $R(0,0,0)$, for MA beneficiaries of the acquirer who are not yet or not ever under a pay-per-patient contract and MA beneficiaries of other insurers, respectively.

	Cohort			
	(1)	(2)	(3)	(4)
	2016	2017	2018	2019
τ_{-4}				-0.109
				(0.037)
τ_{-3}			0.058	0.012
			(0.010)	(0.029)
τ_{-2}		-0.081	0.068	-0.015
		(0.006)	(0.008)	(0.019)
τ_{-1}	0	0	0	0
τ_0	0.239	0.144	0.125	0.231
	(0.033)	(0.007)	(0.008)	(0.020)
τ_1	0.354	0.209	0.074	
	(0.044)	(0.009)	(0.009)	
τ_2	0.327	0.176		
	(0.043)	(0.009)		
τ_3	0.315			
	(0.048)			
Individual FE	X	X	X	X
Year FE	X	X	X	X
Baseline Mean	0.737	0.897	0.794	0.818
Included PCPs	All but ACQ PCP	All but ACQ PCP	All but ACQ PCP	All but ACQ PCP
Treated Observations	4,587	129,170	63,776	8,919
Observations	561,480	561,480	561,480	561,480

Table 4: Differential Coding by Pay-per-Patient Contract Adoption Cohort

Notes: The table presents the event study estimates for the change in patient risk scores arising from pay-per-patient contract adoption across adoption cohort. An observation is a patient-year. We include all PCP practices except for the acquired practice. Across columns, the treated group is made up of MA beneficiaries of the acquiring insurer whose assigned PCP adopted a pay-per-patient contract. The control group is made up of MA beneficiaries of the acquiring insurer who never adopt a pay-per-patient contract and MA beneficiaries of other insurers. These sample requires that we use all diagnoses to compute the risk score. Our estimates are computed using the event study design with individual and time fixed effects from Callaway and Sant'Anna (2021). Control is never treated. Baseline mean is the mean of the treated group at $t = -1$.

noses at their homes, which may increase the likelihood that patients receive "Any Diagnosis" as defined in Equation 6. This may be especially relevant among patients with chronic conditions, such as diabetes, for whom risk scores won't automatically reflect their condition unless physicians re-diagnose them each year. Then, increases on the extensive margin can arise from physicians seeing a patient at all, thereby having any signal of their underlying health status, and from them being more likely to diagnose these patients, as predicted by our framework.

On the intensive margin, PCPs may be more likely to provide more diagnoses than they otherwise would, i.e., they may increase the "Quantity" of diagnoses conditional on having any, as defined in Equation 7.[29,30] Then, increases in the number of diagnoses, conditional on having any diagnosis, can result from physicians becoming more likely to provide the marginal diagnosis and to test patients for potential comorbidities.

Furthermore, physicians may also be more likely to provide diagnoses with higher associated risk scores, or diagnoses with higher "Severity", as defined in Equation 8. Many diagnoses are subjective calls made by providers, and insurers may incentivize them to provide more diagnoses or higher-risk diagnoses. For instance, moving from a no-complication diabetes diagnosis to any complication diabetes diagnosis approximately triples the associated risk.[31]

To evaluate how different margins differentially contribute to the observed increase in risk scores, we first convert our estimates into risk space.[32] Then, We identify these impacts by employing the same empirical strategy as in our main exercise, and conduct event study designs with individual and time fixed effects. Namely, our the treatment group is defined as MA beneficiaries of the acquiring insurer at the acquired PCP and the control group are MA beneficiaries of other insurers at the acquired PCP. We find that the increase in patients' reported risk arises from an increase in both the provision of patients with any diagnosis and the number of diagnoses, conditional on having any diagnosis (Table 5). In contrast, the average severity of diagnoses is decreasing.[33]

Converting our outcomes into risk space is necessary because our three outcomes are in different units, which obfuscates their interpretation. Intuitively, our conversion holds fixed the other margins and imputes what the change in risk score would be if we only turned on one margin at a time. That is, for "Any Diagnosis", we impute what the risk score of the average treated individual with any versus no diagnoses would be in the post periods. For the "Quantity" of diagnoses, we impute what the average risk per diagnosis would be, adjusting for the likelihood of getting any diagnosis. We provide details of how we construct these measures in Appendix C.1.

[29]We define a diagnosis as an HCC with an indicator equal to 1 in the CMS software. In this way, we capture the effective number of diagnoses, but don't count redundant diagnoses, e.g., diagnoses that would not impact the risk score because another diagnosis is equivalent to the software.

[30]In these equations, we denote the number of diagnoses provided to patient i in period t as D_{it}

[31]It is not always physicians who submit the diagnosis codes for a visit, but rather support staff may use physicians' notes to assign the final diagnosis. Then, it may not be that physicians are choosing codes with higher associated severity, but rather the support staff that is taught which code is most advantageous to select when there is more than one option.

[32]We do not need to do this for the average severity, as it is already in risk space.

[33]This decrease is somewhat mechanical as the denominator, the average number of diagnoses, increases with the acquisition.

	(1) Any Diagnosis	(2) Quantity	(3) Severity
τ_{-2}	-0.007	0.025	-0.003
	(0.010)	(0.019)	(0.010)
τ_{-1}	0	0	0
τ_0	0.065	0.045	-0.022
	(0.011)	(0.021)	(0.012)
τ_1	0.140	0.191	-0.023
	(0.013)	(0.027)	(0.012)
τ_2	0.131	0.189	-0.014
	(0.018)	(0.045)	(0.016)
Individual FE	X	X	X
Year FE	X	X	X
Observations	28,204	6,752	6,752
PCP Sample	ACQ PCP	ACQ PCP	ACQ PCP
Insurer Sample	All	All	All

Table 5: Margins of Differential Coding

Notes: The table presents the event study estimates for the change in each margin of differential coding, Any Diagnosis, Quantity, and Severity arising from pay-per-patient contract adoption across adoption cohort. An observation is a patient-year. We restrict attention to patients of the acquired practice, and diagnoses provided by the acquired practice. Across columns, the treated group is made up of MA beneficiaries of the acquiring insurer at the acquired PCP. The control group is made up of MA beneficiaries of other insurers at the acquired PCP. Our estimates are computed using the event study design with individual and time fixed effects from Callaway and Sant'Anna (2021). Estimates are reported in risk space for interpretability, following our conversion procedure as described in Appendix C.1.

4.4 Interpreting Differential Coding

4.4.1 The Acquisition and Pay-Per-Patient Contract Adoption

Our estimates suggest that the acquisition and the adoption of pay-per-patient contracts increase risk-adjustment payments in a similar manner. Indeed, the average impact of the acquisition on total risk of $1,861 per patient per year is not statistically different from the $1,949 impact of pay-per-patient contract adoption.[34] This raises the question of how to interpret the acquisition and the adoption of pay-per-patient contracts together, especially as both of these endogenous treatments are co-timed. Furthermore, because the acquisition is accompanied by the adoption of a pay-per-patient contract at the acquired practice, our estimates capture the bundled treatment effect.

One interpretation of the comparison between our estimates of the bundled treatment at the acquired practice (Table 2 column (2)) and the impact of contract adoption at non-acquired prac-

[34]Note that in order to carry out this comparison we use the estimates that arise from using all diagnoses to construct the risk scores.

tices (Table 3 column (3)) is that it captures the full impact of the acquisition in the absence of contract adoption.[35] This requires two key assumptions. First, that the impacts of acquisition and contract adoption are additively separable, that is, there is no interaction effect.[36] Second, that the acquired practice and the practices that adopt these contracts are comparable.

However, in our setting, we do not believe the latter to be a reasonable assumption, and we therefore do not view these combined exercises as decomposing the individual impacts of acquisition and contracting. Rather, our varied specifications aim to capture the potential heterogeneity in how these incentive alignment mechanisms operate across practices. We discuss the two main reasons behind this interpretation below.

First, there is selection into both acquisition and pay-per-patient contract adoption. Furthermore, if acquisition is needed for contract adoption at the acquired practice but not elsewhere, the selection into each treatment may differ. Indeed, in surveys, over half of physicians in independent practices report the "ease of participation in risk-based payment models" as a key reason for acquisition (Kane, 2025). Then, if the relevant counterfactual for the acquired practice is that, in the absence of acquisition, they would not have adopted a pay-per-patient contract, we would attribute the full impact of the bundled treatment to acquisition.

Second, the impact of contract adoption is fairly heterogeneous across cohorts, suggesting that these contracts may have varying effects on different types of practices. Then, even if the acquired practice had some non-zero probability of adopting one of these contracts in the absence of acquisition, it is unclear what the impact of these contracts at this practice would be. It could be that the likelihood of contract adoption is driven by other practice characteristics which are correlated with higher or lower impacts on diagnoses. However, determining the drivers of contract adoption and their correlation with differential coding is outside the scope of this paper.

4.4.2 Normative Implications of Differential Coding

Our estimates are consistent with Prediction 2 from our conceptual framework, which suggests that the acquisition leads to increased risk-adjustment payments for the insurer in MA. In this section, we evaluate our estimates through the lens of our framework to interpret their normative implications. We further discuss how different extensions to the framework may impact our interpretation.

We start by noting that, as predicted by our model, risk-adjustment payments to the insurer from the government increase, generating a welfare loss of approximately \$$\kappa$1,508 per patient per year in the excess burden of public funds (term *ii*, Equation 4). We then measure the change in diagnostic effort and find that the acquisition led to an increase in diagnostic effort, further decreasing welfare. Lastly, although diagnostic accuracy impacts welfare in our conceptual framework,

[35] This would capture both the direct effect of acquisition as well as its indirect effect through its impact on the likelihood of contract adoption.

[36] Because we don't have variation in contract adoption conditional on acquisition we cannot estimate this interaction effect in our setting.

we cannot measure it and therefore cannot empirically speak to this term.

However, it is worth noting that our conceptual framework does not directly value diagnoses. Furthermore, because our framework is narrowly focused on diagnostic behavior, it does not directly consider the impact on the treatment of these patients. Then, incorporating the relationship between increased diagnostic provision and treatment allows for potential improvements in patient welfare. This is because, if patients value receiving therapeutic care, then increases in treatment that accompany increased diagnoses could prove to be beneficial for welfare.[37]

We address this possibility in two ways. First, in Appendix B, we extend the model to include physicians' value from providing therapeutic care, in addition to their diagnostic care. Physicians receive utility both from the value of providing care for unhealthy patients as well as from the fee-for-service payments they receive from the insurer to provide such care. For the acquisition to increase the utility of care, it would have to be the case that both the acquisition increases treatment and that physicians were previously providing less care than would be preferred by the social planner. Furthermore, the latter contradicts the conventional concern regarding fee-for-service payments for physician procedures, which holds that these payment arrangements incentivize the over-provision of care.

Our modified framework suggests that a commensurate increase in treatment does not accompany the increase in diagnoses. Intuitively, in our modified framework, as providing treatment is costly and more beneficial for unhealthy than healthy patients, the PCP only wants to treat those patients it is fairly certain to be unhealthy. The PCP, therefore, does not treat any undiagnosed patients. For diagnosed patients for whom the PCP does not exert effort, they only treat those with a lower probability of being healthy, and the acquisition shifts this cutoff so that the PCP is stricter and treats fewer patients. The acquisition further shifts the cutoff for exerting effort and treating those patients who are revealed to be unhealthy, leading to a decrease in the number of treated patients near the lower bound. Nevertheless, as the acquisition increases the diagnosis and treatment of patients who appear to be healthier, these patients may receive improved care. All in all, however, our framework predicts that the increase in diagnoses is strictly larger than the change in treatment, even when the latter increases.

Second, we conduct an empirical exercise to gauge whether this prediction is correct - does an increase in treatment accompany the increase in diagnostic effort? First, we measure the number of procedures received by each patient each year, as well as the total number of individual diagnoses. Then, we classify medical procedures into "Therapeutic" procedures and "Diagnostic" procedures,where the former capture treatment effort and the latter capture diagnostic effort. Lastly, we conduct an event study design to estimate the impact of the acquisition on each of these outcomes.

As before, we restrict attention to patients of the acquired practice, define our treated group

[37]Note that for increases in treatment to be beneficial it would have to be the case that physicians were previously under-providing treatment.

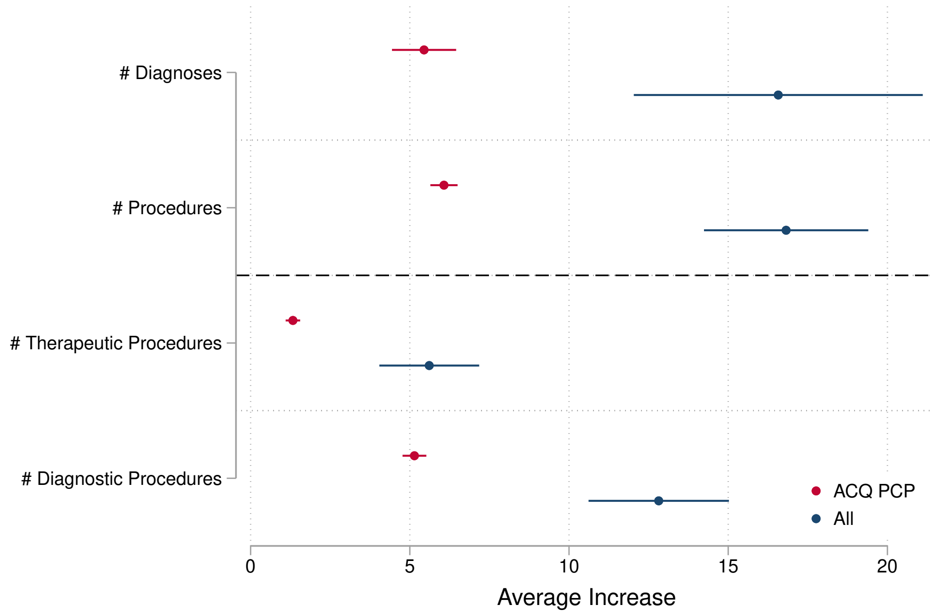

Figure 4: Heterogeneity of Impacts across Types of Procedures

Notes: The figure presents the event study estimates for the change in the number of diagnoses and procedures, as well as a break down along whether procedures are classified as therapeutic or diagnostic. An observation is a patient-year. We restrict the sample to our Acquired PCP sample, that is, we keep only patients of the acquired practice. The treated group are MA beneficiaries of the acquiring insurer at the acquired practice. The control group are MA beneficiaries of other insurers at the acquired practice. The estimates differ in the provenance of the diagnoses and procedures used for the outcome. The ACQ PCP estimates restrict attention to diagnoses and procedures carried out by the acquired PCP. The All estimates use all diagnoses and procedures received by these patients at any physician. Our estimates are computed using the event study design with individual and time fixed effects from Callaway and Sant'Anna (2021).

as the MA beneficiaries of the acquiring insurer at the acquired practice, and our control group as the MA beneficiaries of other insurers at the acquired practice. We conduct the exercise using only those procedures and diagnoses provided to these patients by the acquired practice as well as those provided by any practice. We present our results in Figure 4, where we find that, although diagnoses and total procedures present similar increases, the vast majority of these procedures are diagnostic rather than therapeutic.[38] Then, in practice, the increase in diagnoses is at least not fully met by a corresponding increase in treatment. This divergence suggests that the increase in diagnoses is not likely to be driven by an increase in the need for patient care.

5 Managing the Cost of Specialist Care

In this section, we empirically evaluate whether the acquisition delivers cost savings in specialist care through referral steering, as illustrated by our framework. We find that the acquired physi-

[38]Note that because we include individual and time fixed effects, these estimates are not raw averages and thus do not need to sum up to the total impact on procedures.

cian practice becomes more cost sensitive and steers referrals towards more cost-effective specialists for Commercial beneficiaries of the acquiring insurer. To do this, we start by constructing our referral sample. Then, we present some descriptive findings on referral patterns, before writing and estimating a model of PCP referrals. Lastly, we discuss the interpretation of our findings.

5.1 Referral Sample

We construct our sample by first selecting all specialist claims and then restricting them to those that may feasibly result from a referral. Then, we identify the referring PCP for each claim before restricting to our final sample.

We consider two broad categories of specialist services: inpatient claims (IP), such as hospital admissions, and outpatient claims (OP), such as professional services. We classify inpatient claims using the diagnosis-related-group (DRG) classification, as is standard in the literature (Cutler et al., 2001). DRGs group together patients with similar diagnoses and care needs, accounting for the patient's condition, the complexity of the procedure, and any complications.[39] Outpatient claims are classified via Berenson-Eggers Type of Service codes (BETOS), as described in Berenson and Braid-Forbes (2020). BETOS codes collapse granular billing codes into clinically meaningful categories, like office visits or tests, therefore facilitating the analysis of each category of care.[40]

For each type of specialist service, we first drop those that could not feasibly be the result of a referral. We do this by dropping claims labeled as Emergency, Urgent, or Newborn. We further restrict to valid claims and specialists with at least 20 claims per year. We conduct this exercise separately for inpatient and outpatient claims. Then, we construct a referring PCP indicator for each claim. We define a patient as being referred by their PCP if, in the year of their inpatient claim, they had previously visited their assigned PCP. This referral assignment process follows that of Brot-Goldberg and de Vaan (2018).

Lastly, we separate our sample into Commercial and MA beneficiaries, beneficiaries of the acquiring insurer versus any other insurer, and patients of the acquired practice versus any other practice. Furthermore, we flag those beneficiaries of the acquiring insurer who are covered under a pay-per-patient contract at their assigned practice, as before.

5.2 Descriptive Findings

Because patients often rely on referrals to see a specialist, primary care physicians can significantly influence whether patients visit a specialist and which specialist they choose. Because the PCP only affects specialist spending indirectly through their referral choice, the level and variation in specialist costs matter for the savings that can be achieved through a different referral choice. In this section, we present some descriptive results on specialist price variation and on our referral sample.

[39]For instance, a patient undergoing a knee replacement may be classified under DRG 469 – Major joint replacement or reattachment of lower extremity with MCC (Major Complication or Comorbidity)

[40]For instance, code M5C is assigned for an ophthalmology specialist office visit.

The scope for specialist savings depends on the cost of specialist care and its variation. In Appendix Table F.3, we compare specialist prices for inpatient and outpatient claims across market segments. We find that, for both types of specialist care, both the level and the variance are higher in the Commercial segment than in the MA segment. This is consistent with prior literature showing that MA plans anchor prices to Traditional Medicare rates, mechanically lowering price variation (Trish et al., 2017). The Commercial segment, in contrast, lacks this anchor and consequently exhibits both a higher mean and variance in specialist prices. Although this pattern holds both for inpatient and outpatient claims, it is further worth highlighting that, although outpatient claims are a lot more common, their costs per claim are significantly smaller than those of inpatient claims. These patterns highlight that there is a larger scope for savings in the Commercial segment than in the MA segment.

In Figures 5 and 6, we plot the total number of inpatient and outpatient referrals, respectively. We further group patients by their market segment (MA or Commercial), their insurer (acquirer or all others), and their primary care physician (acquired practice or any other practice). As implied by our conceptual framework, as more MA beneficiaries are diagnosed and become eligible for referral, the number of referrals among MA patients increases. However, there is no differential impact in the Commercial segment, where the number of referrals trends down for all patients. This pattern holds true both for inpatient and outpatient claims.

Examining these figures highlights a first concern when evaluating the impact of the acquisition on inpatient claims among MA beneficiaries. Namely, our sample includes only two referrals for MA beneficiaries of other insurers during the pre-acquisition period.[41] Because of this, we cannot evaluate the impact of the acquisition on inpatient referrals among MA patients of the acquired practice. Furthermore, although our framework does not consider it, the physician might increasingly refer their patients to specialists who are also part of the integrated firm. However, the acquiring insurer does not own any hospitals in our period, alleviating this concern among inpatient claims. Furthermore, we show in Appendix Figure E.3, that there is no change in the share of within-group referrals for outpatient claims, that is, we find no evidence of increased self-steering.

Lastly, our first model prediction (1) illustrates how integration decreases the "intensity" of physicians' referrals. That is, the physician becomes more sensitive to the prices of specialists, steering patients towards more affordable options. Not only is this what our detailed price data is best suited to analyze, but in the Commercial segment, where there is a larger scope for savings, we find no indication of changes on the number of referrals. Note that this is not surprising in the context of private insurance, as price variation is responsible for an estimated 70% of geographic variation in healthcare spending in this context (Newhouse and Garber, 2013). Then, we focus on steering towards more cost-effective specialists for the rest of the section.

[41] This is because, although the acquired practice cares for 1119 beneficiaries of other insurers in 2016, inpatient referrals are fairly uncommon. Indeed, the referral rate at the acquired practice is smaller than 1% by group, as defined by the interaction of segment and insurer.

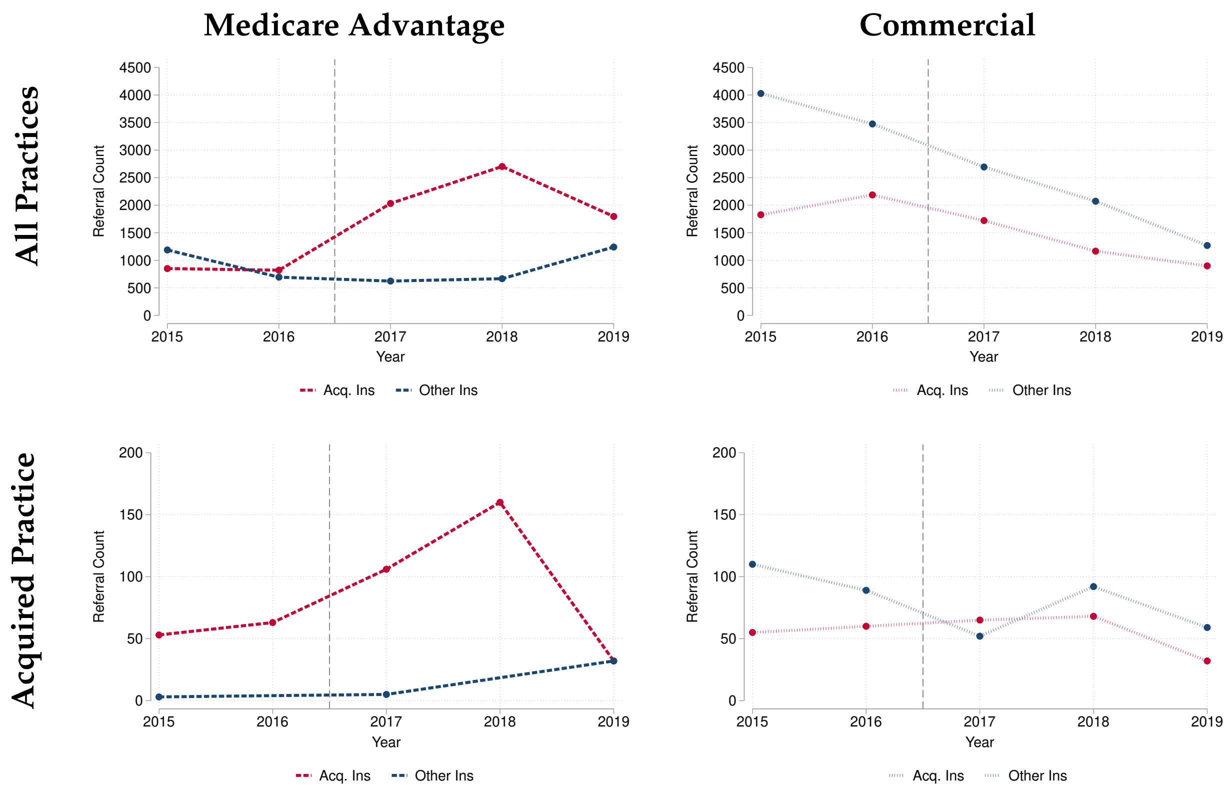

Figure 5: Trends in Inpatient Referrals

Notes: The figures illustrate how the number of inpatient referrals for beneficiaries of the acquiring insurer and other insurers across market segments MA evolve over the course of our sample, both for patients at the acquired practice and elsewhere. We consider inpatient claims, i.e., hospital admissions, classified by DRGs. We then restrict to those claims that may results from a referral and identify the referring PCP. Then, we plot the number of referrals made by the acquired practice in the bottom row, and by all practices in the top row. The vertical line in each panel indicates the 2017 acquisition date.

To evaluate the impact of the acquisition on steering towards more cost-effective specialists, one may be tempted to compare the realized prices of referrals before and after the acquisition. However, the realized price of a medical visit is not the object of interest. Rather, when making referral decisions, physicians care about the expected price of the specialist. Furthermore, although the realized price may act as a proxy for the object of interest, it is a fundamentally noisy one, and there are unobservables that we cannot control for to correct for this noise.

5.3 Model Set Up

We evaluate the impact of integration on the cost of specialist care by estimating a specialist cost index and a PCP-referral choice model. In this way, we overcome the noisiness of realized prices. Furthermore, this model allows us to simulate and evaluate counterfactual referrals in the absence of acquisition.

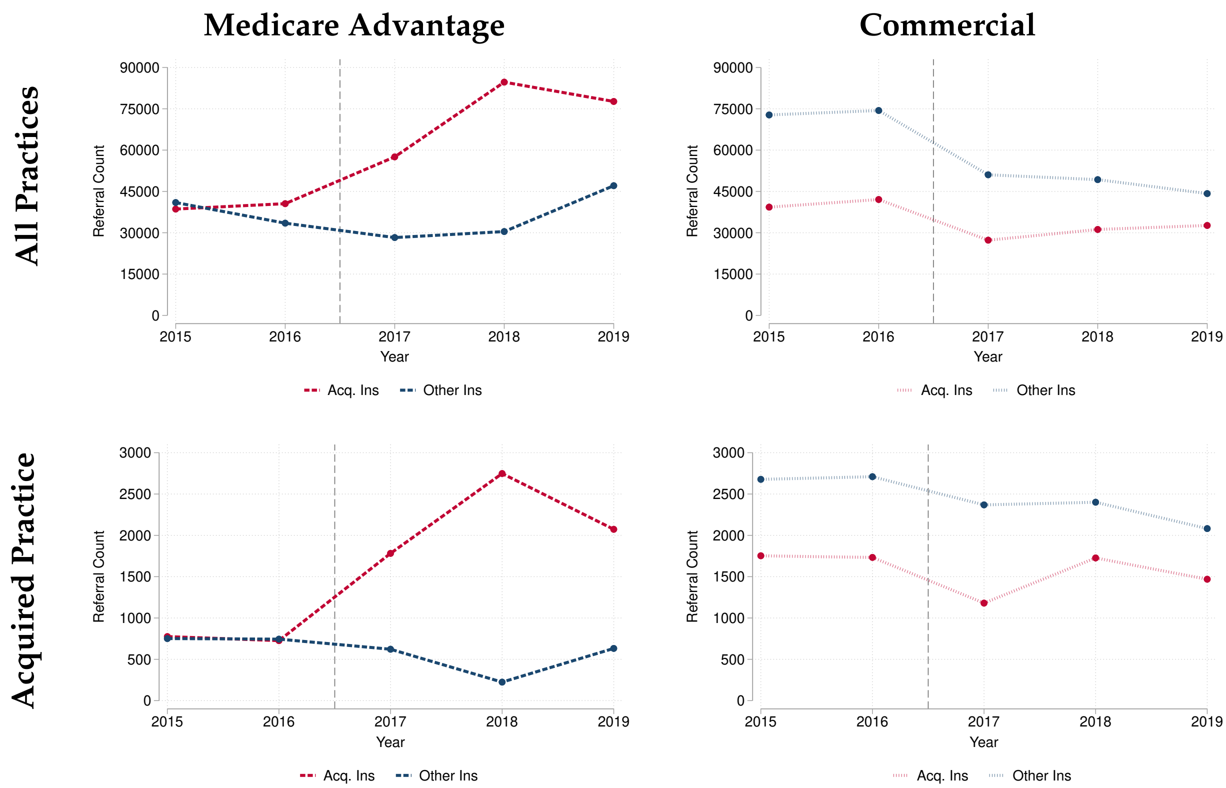

Figure 6: Trends in Outpatient Referrals

Notes: The figures illustrate how the number of outpatient referrals for beneficiaries of the acquiring insurer and other insurers across market segments MA evolve over the course of our sample, both for patients at the acquired practice and elsewhere. We consider outpatient visits, as classified by BETOS codes. We restrict to those claims that may results from a referral and identify the referring PCP. Then, we plot the number of referrals made by the acquired practice in the bottom row, and by all practices in the top row. The vertical line in each panel indicates the 2017 acquisition date.

Demand Model. To study the impact of the acquisition on the steering of patients towards cheaper specialists, we estimate a demand model of joint PCP-patient specialist choices, as in Ho and Pakes (2014a). Let the referring PCP k and their patient i, covered under insurance plan $g(i)$ jointly choose which specialist j to be referred to for care in period t. Note that, as we condition on referrals, we do not include an outside option. Then, we define the referral decision problem as follows:

$$u_{kijt} = \eta_{k,g(i),t} \text{ Expected Specialist Cost}_{g(i),j,t} + \beta \text{ dist}_{ijt} + \delta_j + \nu_{kijt} \tag{9}$$

$$\eta_{k,g(i)t} = \bar{\eta} + \eta^{\text{Acq Ins \& Post}}_{k,g(i),t} \mathbb{1}_{g(i)=\text{Acq Ins \& } t=\text{Post}} + \eta^{\text{Acq Ins}}_{k,g(i),t} \mathbb{1}_{g(i)=\text{Acq Ins}} + \eta^{\text{Post}}_{k,g(i),t} \mathbb{1}_{t=\text{Post}} \tag{10}$$

In our specification, specialists j are vertically differentiated by quality δ_j and horizontally differentiated by their distance to each patient dist_{ijt}. Furthermore, specialists are characterized by an expected cost index for each insurance plan each period, for which the referring physician has cost sensitivity $\eta_{k,g(i),t}$, which we allow to vary by the patient's insurance plan and period. Lastly, $\nu_{kijt} \sim T1EV$ is an idiosyncratic shock to preferences for specialists by PCP-patient-specialist-period. Then, when choosing which specialist to refer their patient to, the PCP trades off the expected cost of care with specialist quality and distance.

To capture the impact of integration on cost sensitivity, we impose a two-way fixed effects structure on the cost sensitivity parameter $\eta_{k,g(i),t}$. We interpret the coefficient $\eta^{\text{Acq Ins \& Post}}_{k,g(i),t}$ as capturing the impact of the acquisition on the referring physician's price responsiveness.[42] In this way, our estimation strategy reduces to a differences-in-differences approach in η space. The identifying assumption is that, absent the acquisition, the acquired practice's price responsiveness would have evolved in a parallel fashion for beneficiaries of either insurer. Note that, if we don't include the specialist random effect δ_j to control for quality, as we do in some specifications, the differences-in-differences assumption can be strengthened to maintain the same interpretation of the coefficient of interest. Namely, if the η parameter confounds preferences over price and quality, and this confounding evolves in the same way across groups, then the parallel trends assumption differentiates out this confounding.

Cost Index Construction. To estimate our demand function, we first construct a measure of the expected cost of referring a patient to each specialist. To construct this cost index, we use all claims at each specialist, grouped by the type of specialist care being provided, insurance plan, and period. We further carry out two corrections to our cost index: a closed-form leave-one-out correction and a hierarchical empirical Bayes shrinkage procedure to reduce noise.

We define an observation as a claim for patient i, under insurance plan $g(i)$, at specialist j, in period t, for a specialist service group s. We divide specialist services into two categories, inpatient and outpatient services, and then into groups within each category. For inpatient claims, groups are defined at the diagnosis-related-group (DRG) claim classification.[43] For outpatient claims, groups are defined by their BETOS code classification.[44] Then, within each group s, plan $g(i)$, and period t we retrieve the specialist cost index from the following linear regression:

$$\$\text{ Amount}_{ijt} = \alpha_{g(i),s,t} + \text{Spe FE}_{g(i),j,s,t} + \mathbf{X}_{ijts} + \varepsilon_{ijts} \tag{11}$$

where $\mathbf{X}$ includes year, gender, pay-per-patient contract, and MA fixed effects, as well as controls for patient risk score. Then, the specialist cost index varies at the specialist service, period, and insurance plan level. We further demean it so that it is centered around zero. We then follow

[42] Note that this is equivalent to estimating the elasticities per group and then estimating a differences in differences model, as there is staggered treatment adoption.

[43] DRGs group together patients with similar diagnoses and care needs, accounting for the patient's condition, the complexity of the procedure, and any complications.

[44] BETOS codes collapse granular billing codes into clinically meaningful categories, like office visits or tests, therefore facilitating the analysis of each category of care.

Miller (1974) to perform a post-estimation, closed-form leave-one-out correction, and carry out a two-dimensional hierarchical empirical Bayes shrinkage procedure to reduce noise. We shrink our estimates within plans and across specialists and within specialists across plans to leverage both plan and specialist-level information. We detail these correction procedures in Appendix D.1.

To construct our cost index, we utilize all the claims for each specialist that could feasibly result from a referral, as defined in Section 5.1. We conduct this exercise separately for inpatient and outpatient claims, and construct the cost index at the specialist service group as defined above. We express dollar amounts in thousands of dollars, and use patients' pre-acquisition risk score as a control.[45]

5.4 Results

We find that the acquisition led to an increase in the price responsiveness of the PCP when referring Commercial patients for inpatient care with specialists, resulting in savings of approximately \$300 per inpatient and \$26 per outpatient referral. However, when evaluating the impact on outpatient referrals for MA beneficiaries, we find an average increase in specialist costs of \$233 per outpatient referral. Nevertheless, these estimates are noisy and not distinguishable from zero for a number of outpatient specialist categories.

5.4.1 The Impact Vertical Integration on Referral Steering

To estimate our model, we first construct the specialist cost index. Then, we present descriptives on how the average cost index of the acquired practice's referrals evolve over time. Lastly, we estimate our demand model and the implied referral savings.

We first construct our specialist cost index as described in Section 5.3. In Appendix Table F.4 we present summary statistics for the inpatient and outpatient cost indices in our sample. Here, we see that, even after shrinking, there is still substantial variation in our expected cost measure. To illustrate the impact of our shrinkage procedure on our cost index, in Appendix Table F.5, we present summary statistics for the inpatient cost index before and after shrinkage.

Then, to evaluate how the acquisition impacts referral steering at the acquired practice, we start by inspecting trends in the average referral cost index over time. This further helps build intuition for the variation that identifies our model estimation. To do this, we restrict attention to inpatient referrals of Commercial patients of the acquired practice. As in our event study designs, we compare beneficiaries of the acquiring insurer to those of other insurers. In Figure 7, we present the average cost index of the realized referrals for these patients. This figure illustrates the variation in expected cost that our model estimates are meant to capture. Based on the trend of decreased average cost for referrals, we expect the acquired practice to steer Commercial beneficiaries of the acquiring insurer towards more cost effective specialists. However, these descriptive

[45]We don't use the risk scores from the post period given the evidence that the acquisition impacts these, i.e., they are not invariant to the treatment. We also utilize the CMS software to calculate risk scores for the Commercial segment, which we will use in this exercise.

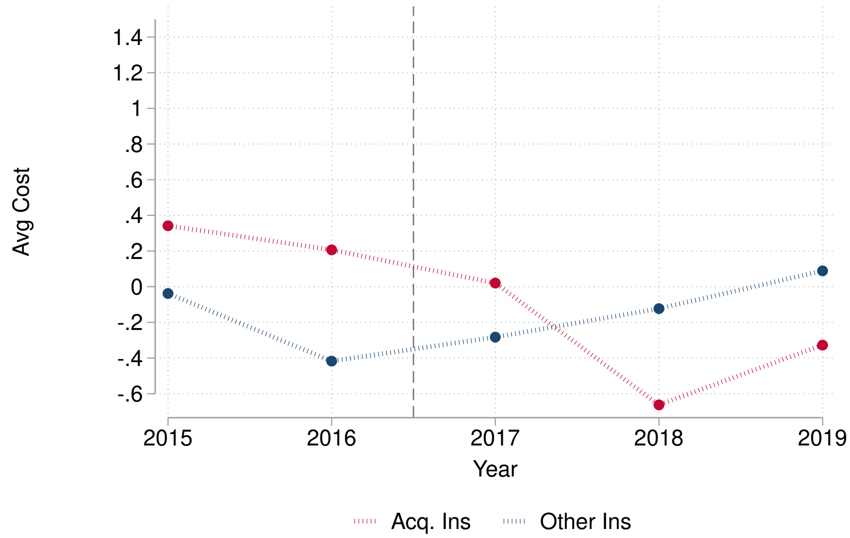

Figure 7: Cost Index for Commercial IP Claims over Time

Notes: The figure illustrates how the average cost index for realized inpatient referrals for Commercial patients at the acquired practice evolve over time. We restrict referrals to those outpatient claims that may results from a referral and identify the referring PCP. An observation is a referral. We restrict our sample to Commercial patients of the acquired practice and compare beneficiaries of the acquiring insurer to those of other insurers. The vertical line in each panel indicates the 2017 acquisition date.

averages may not capture other aspects of the referral decision, such as distance between patients and specialists, or the choice set available to the PCP when providing referrals.

Next, we estimate our demand model to determine the impact of the acquisition on the physician's price sensitivity, and use our estimates to conduct counterfactual simulations and convert the changes in price sensitivity into estimated savings. To estimate cost savings, we simulate referrals under the estimated model and under a counterfactual where the acquisition did not occur, i.e., $\hat{\eta}_{it}^{\text{Acq Ins \& Post}} = 0$. We then compare the expected costs under each model and interpret the difference as the total savings arising from the change in price sensitivity caused by the acquisition.

We estimate our model for a variety of samples and under two different specifications. We first estimate the model for inpatient specialist visits, restricting to only Commercial beneficiaries first, and then including all beneficiaries.[46] Then, we conduct estimate the model for outpatient services as classified by their BETOS codes, restricting to only Commercial beneficiaries first, and then restricting to only MA beneficiaries. For each of these samples, we conduct the estimation with and without including the random effect δ_j.

We interpret the differences in the specification that includes or does not include the random effect as follows. As discussed in Section 5.3, the differences-in-differences structure of the cost sensitivity parameter can difference out the impact of preferences over quality under certain as-

[46]We do not conduct this exercise for MA separately because, as noted in Section 5.2, we do not have a large enough sample in MA.

sumptions. Namely, if the estimates of cost sensitivity for each group face the same confounding from quality, and treatment does not impact preferences over quality, then our structure captures solely changes in cost sensitivity. Including the random effect is meant to capture the unobserved quality directly, avoiding confounding with quality in our estimates of η. Nevertheless, because the random effect is constant over time, it still requires that we assume it is not impacted by treatment.

We present our model estimates for inpatient referrals in Table 6. We find that, after integration, the PCP is more price sensitive for referrals of Commercial beneficiaries. That is, the estimated coefficient $\hat{\eta}_{kg(i)t}^{\text{Acq Ins \& Post}}$ is negative in columns (1) and (2). Furthermore, converting our estimated price sensitivity into cost savings yields savings of \$255-\$321 per referral, impacting 165 referrals for Commercial beneficiaries of the acquiring insurer. Furthermore, comparing our estimates with and without the random effect δ_j reveals that the coefficient of interest is not significantly different across specifications, with the caveat that the standard error for the random coefficient is very noisily estimated. Lastly, including the random effect does impact the coefficient on distance. This is consistent with the notion that quality and willingness to travel may be correlated.[47]

We conclude our evaluation of the impact of integration on inpatient referral steering by our model estimates among all beneficiaries in Table 6 columns (3) and (4) yields estimates that are not distinguishable from zero. We interpret this variation in estimates as arising from noise in the MA segment, largely due to the small number of referrals, which also prevents us from analyzing this segment separately. Furthermore, because of how MA prices are negotiated, we expect less variation in this segment as well, which could lead to poor identification of the parameters of interest. Then, although we cannot draw conclusions on the impact on MA beneficiaries from this, it is consistent with the idea that the acquisition has a different impact on MA beneficiaries than it does on Commercial beneficiaries.

Next, we shift our focus to steering for outpatient specialist procedures. This sample expansion allows us to overcome the sample size limitations of inpatient referrals, and evaluate the impact for Commercial and MA beneficiaries separately. We present a summary of our estimates for Commercial beneficiaries in Table 7 and for MA beneficiaries in Table 8. This summary is comprised of the estimates for the coefficient of interest for seven categories of outpatient procedures, including home visits, specialist office visits, and major cardiovascular and orthopedic procedures. The full estimates are presented in Appendix Tables F.6, F.7, F.8, and F.9. These estimates highlight that, although there is heterogeneity across types of outpatient specialist care, overall, the PCP steers Commercial patients towards cheaper specialist care, yielding average savings of \$26 per referral. However, among MA patients, there are no such savings. Most of our estimates are not distinguishable from zero, and those that are distinguishable yield cost increases rather

[47] Indeed, Ho and Pakes (2014a) find that under pay-per-patient contracts, physicians will hold quality constant while trading off price and distance.

	Commercial		All Segments	
	(1)	(2)	(3)	(4)
$\hat{\eta}^{\text{Acq Ins \& Post}}_{kg(i)t}$	-0.704	-0.923	-0.084	-0.425
	(0.324)	(0.524)	(0.195)	(0.355)
$\hat{\eta}^{\text{Acq Ins}}_{kg(i)t}$	0.879	0.773	0.789	0.250
	(0.219)	(0.431)	(0.164)	(0.215)
$\hat{\eta}^{\text{Post}}_{kg(i)t}$	0.086	0.331	0.140	0.236
	(0.267)	(0.369)	(0.165)	(0.262)
$\hat{\bar{\eta}}$	-0.301	-0.376	-0.250	-0.221
	(0.189)	(1.001)	(0.136)	(0.191)
$\hat{\beta}$	0.566	3.568	0.237	1.832
	(0.120)	(0.558)	(0.118)	(0.299)
σ_δ		20.114		19.977
		(115.023)		(44.869)
Observations	682	682	1136	1136
Random Effect	-	X	-	X
Referring PCP	ACQ PCP Only	ACQ PCP Only	ACQ PCP Only	ACQ PCP Only
Insurers	All	All	All	All
Δ Cost ($ per ref)	-255	-321	-43	-253
N Impacted Referrals	165	165	463	463

Table 6: Integration and Inpatient Referral Steering

Notes: The table presents the model estimates for the change in the acquired PCP's cost sensitivity for inpatient referrals across specifications and samples. An observation is a referral. We restrict our sample to our Acquired PCP sample, that is, we keep only patients of the acquired practice. In columns (1) and (2), the treated group are Commercial beneficiaries of the acquiring insurer at the acquired practice. The control group are Commercial beneficiaries of other insurers at the acquired practice. Columns (3) and (4) evaluate the impact across all beneficiaries, Commercial and MA. Then, the treated group are all beneficiaries of the acquiring insurer at the acquired practice and the control group are all beneficiaries of other insurers at the acquired practice. In columns (1) and (3) we do not include random effects. In columns (2) and (4) we do include random effects.

than savings. Then, our estimates imply an average cost increase of $234 per referral.

Lastly, we discuss some concerns with our model estimates. A first potential concern with our estimates arises from the fact that we recompute the cost index for both the pre- and post-periods. We do this to ensure our index captures the fact that the acquisition may impact both negotiated prices and the provider network available to the insurer's beneficiaries. However, this may impact our estimates. Whereas our estimates would not be affected by the insurer negotiating cheaper prices across the board, they would be biased towards no impact if the price distribution became more compressed. Conversely, if we thought that the insurer may be more likely to include more expensive providers in their network because they can prevent physicians from referring patients to them, our estimates would overstate the increase in price sensitivity.

	$\hat{\eta}^{\text{PC \& Post}}_{kg(i)t}$	$Std.Err$	Obs	Rand Eff.	Δ Cost ($ per ref)	Referrals
Home Visits	-0.757	0.123	754	-	-2112	289
	-0.636	0.218	754	X	-1033	289
Specialist Visits	0.117	0.049	10497	-	120	1854
	0.236	0.061	10497	X	156	1854
Major Proc	0.211	0.091	2924	-	185	656
	0.246	0.096	2924	X	189	656
Major Cardio Proc	-0.253	0.116	404	-	-259	127
	0.334	0.435	404	X	244	127
Major Ortho Proc	0.271	0.046	2310	-	235	576
	-0.005	0.100	2310	X	-5	576
Eye Proc	-0.058	0.035	1484	-	-104	304
	0.056	0.096	1484	X	75	304
Ambulatory Proc	0.064	0.056	4956	-	58	1150
	0.003	1.209	4956	X	3	1150

Table 7: Integration and Commercial Outpatient Referral Steering

Notes: The table presents the model estimates for the change in the acquired PCP's cost sensitivity for outpatient referrals across specifications and types of outpatient procedures. An observation is a referral. We restrict our sample to our Acquired PCP sample, that is, we keep only patients of the acquired practice. We further restrict our sample to only include Commercial beneficiaries. The treated group are Commercial beneficiaries of the acquiring insurer at the acquired practice. The control group are Commercial beneficiaries of other insurers at the acquired practice. For each type of outpatient procedure, we estimate the model with and without random effects.

A second concern arises if we think that the change in cost is not the result of increased cost sensitivity, but rather of steering within the integrated firm. This is not a concern among inpatient claims because the insurer does not own any hospitals in the market in this period, and therefore, the steering we find is not driven by ownership. Furthermore, we do not find an increase in the share of outpatient referrals within the integrated firm (Appendix Figure E.3). Lastly, the referral claims are in-network, supporting the idea that the PCP is choosing among a patient's choice set as defined by their provider network.

5.4.2 The Impact of Pay-per-Patient Contract Adoption

Next, we conduct an exercise to evaluate whether, in the absence of an acquisition, pay-per-patient contracts may suffice to induce such an increase in PCPs' price responsiveness and find that this is not the case. To do this, we drop referrals provided by the acquired PCP and define treatment as the adoption of a pay-per-patient contract. Furthermore, because contract adoption occurs almost exclusively for MA beneficiaries of the acquiring insurer, we restrict our sample to referrals of MA beneficiaries only. Lastly, due to the staggered adoption of these contracts, our two-way fixed

	$\hat{\eta}^{\text{PC \& Post}}_{kg(i)t}$	$Std.Err$	Obs	Rand Eff.	Δ Cost ($ per ref)	Referrals
Home Visits	0.532	0.080	3040	-	450	2552
	0.996	0.176	3040	X	784	2552
Specialist Visits	0.195	0.026	8182	-	231	5044
	0.197	0.043	8182	X	227	5044
Major Proc	0.270	0.115	955	-	243	639
	0.548	1.580	955	X	511	639
Major Cardio Proc	-0.151	0.275	428	-	-161	287
	-0.139	0.436	428	X	-138	287
Major Ortho Proc	-0.140	0.141	1168	-	-138	761
	0.366	0.123	1168	X	369	761
Eye Proc	0.266	0.032	2312	-	153	1003
	-0.014	0.030	2312	X	-10	1003
Ambulatory Proc	0.195	0.151	2086	-	184	1371
	0.036	0.138	2086	X	38	1371

Table 8: Integration and MA Outpatient Referral Steering

Notes: The table presents the model estimates for the change in the acquired PCP's cost sensitivity for outpatient referrals across specifications and types of outpatient procedures. An observation is a referral. We restrict our sample to our Acquired PCP sample, that is, we keep only patients of the acquired practice. We further restrict our sample to only include MA beneficiaries. The treated group are MA beneficiaries of the acquiring insurer at the acquired practice. The control group are MA beneficiaries of other insurers at the acquired practice. For each type of outpatient procedure, we estimate the model with and without random effects.

effects setup does not perfectly translate to this exercise. We therefore estimate our coefficients for the price sensitivity on inpatient visits by cohort instead, and present our results in Table 9.

We find that PCPs that adopt these contracts exhibit a decrease in price responsiveness after adoption, although not all our estimates are statistically distinguishable from zero. It is worth noting, however, that because we compare across physicians in this specification, these estimates could capture how the selection into contract adoption correlates with specialist costs rather than their impact. For instance, if PCPs that are associated with a large system are more likely to adopt these contracts and PCPs in large health systems are more likely to refer within their system, then their referrals may capture the average cost of specialist care within these health systems.

5.4.3 Interpretation

These varied impacts across groups suggest that incentive alignment alone may not be enough to generate referral steering. At first glance, our findings suggest that contracting does not achieve the desired referral steering, but the acquisition does, at least among Commercial beneficiaries. This is consistent with the idea that aligning incentives may not be sufficient for physicians to

	Cohort					
	2017		2018		2019	
	(1)	(2)	(3)	(4)	(5)	(6)
$\hat{\eta}^{\text{PC \& Post}}_{kg(i)t}$	0.136	-0.008	0.374	0.449	0.616	0.130
	(0.077)	(0.109)	(0.076)	(0.192)	(0.152)	(0.329)
$\hat{\eta}^{\text{PC}}_{kg(i)t}$	0.443	0.300	0.397	0.117	0.261	0.190
	(0.072)	(0.094)	(0.090)	(0.207)	(0.123)	(0.285)
$\hat{\eta}^{\text{Post}}_{kg(i)t}$	-0.006	-0.251	-0.024	-0.150	-0.046	-0.143
	(0.046)	(0.236)	(0.032)	(0.147)	(0.047)	(0.264)
$\hat{\bar{\eta}}$	-0.075	-0.062	-0.070	-0.113	-0.071	-0.122
	(0.043)	(0.668)	(0.026)	(1.581)	(0.027)	(1.509)
$\hat{\beta}$	-0.109	-0.301	-0.211	-0.275	-0.237	-0.348
	(0.019)	(0.069)	(0.026)	(0.072)	(0.034)	(0.088)
σ_δ		3.273		3.285		3.254
		(154.8)		(277.0)		(291.3)
Observations	39798	39798	37522	37522	36608	36608
Random Effect	-	X	-	X	-	X
Δ Cost ($ per ref)	97	-6	234	292	503	137
Referrals	2670	2670	747	747	100	100

Table 9: Pay-per-Patient Contracts and MA Inpatient Referral Steering

Notes: The table presents the model estimates for the change in PCP's cost sensitivity for inpatient referrals after adopting a pay-per-patient contract. We conduct out estimation by adoption cohort to overcome concerns with staggered treatment adoption. An observation is a referral. We restrict our sample to MA beneficiaries because they are the only population that adopts pay-per-patient contracts. We drop patients of the acquired practice. In each column, the treated group are MA beneficiaries of the acquiring insurer whose PCP adopts a pay-per-patient contract in a given year. The control group are MA beneficiaries of other insurers and MA beneficiaries of the acquiring insurer who never adopt a pay-per-patient contract. In columns (1), (3), and (5) we do not include random effects. In columns (2), (4), and (6) we do include random effects.

steer referrals towards more cost-effective specialists. Instead, to effectively steer patients towards cheaper specialists, PCPs need not only the right incentives but also the right information. For instance, Cho (2025) finds that integrated PCPs adopt information technologies that may facilitate effective referrals, such as EHR systems. In this way, integration may be a necessary condition for overcoming information asymmetries and achieving effective steering.

However, because these contracts are only adopted in the MA segment, our findings across treatments are not directly comparable. We can therefore not rule out the possibility that the underlying difference in steering ability is related to the market segment rather than the mechanism for incentive alignment. These differences across segments align with the expected variation arising from cost differences. Furthermore, our findings indicate that the acquisition increases the average cost of outpatient referrals among MA beneficiaries. Lastly, because we don't observe contract adoption for Commercial beneficiaries, we cannot directly speak to whether pay-per-patient contracts would be effective for referral steering in our setting.

5.5 Other Impacts on Beneficiaries' Cost of Care

ER Use. A commonly cited advantage of integration is the reduction of unnecessary or preventable emergency care. Through their provision of primary and preventive care, PCPs can impact the ER use of their patients. For instance, increased after-hours and telephone availability can lead patients to substitute ER use for PCP care. Furthermore, through preventive care, PCPs may decrease the need for their patients to visit the ER by preventing such episodes, for example, by managing patients with diabetes' A1C levels. We find that, overall, the amount of spending on ER visits slightly increases for patients of the acquiring insurer compared to other patients at the acquired practice. Nevertheless, we leverage the Billings algorithm (Billings et al., 2000) to evaluate the composition of ER care. We find a slight shift in the composition of this care, towards addressing necessary and hard-to-prevent emergencies for MA beneficiaries. We present these results in Appendix Table F.10.

Steering and Price Discrimination As briefly discussed in Section 2.3, the acquisition does not imply exclusivity. That is, the acquired PCP still cares for beneficiaries of other insurers after the acquisition. Nevertheless, there remains the possibility that the acquiring insurer steers its beneficiaries to the acquired PCP. Indeed, we find that an increase in the likelihood that the insurer's MA beneficiaries are patients of the acquired PCP after the acquisition. However, this is also true for Commercial beneficiaries of other insurers. We present these results in Appendix Table F.11.

One interpretation for steering beneficiaries towards the acquired practice is that it provides an opportunity to better manage these patients' care and its associated costs. For MA beneficiaries specifically, this can provide the integrated firm more control over its beneficiaries' diagnoses. As for the increase in Commercial beneficiaries of other insurers, this may be beneficial for the integrated firms' revenues, as other insurers are directing more expenditures towards the integrated firm.

6 Discussion

Overall Implications of the Acquisition. We start by focusing on the overall impact of the acquisition in the Medicare Advantage segment. In Medicare Advantage, our findings are consistent with model prediction 2, as the acquired practice increases its diagnostic provision for patients of the acquiring insurer by an average of \$1508 per patient per year. However, we do not find evidence supporting model prediction 1 among MA beneficiaries. Instead, we find that the acquired practice provides more referrals for the MA beneficiaries of the acquiring insurer, and they do not become more price sensitive when providing these referrals. Furthermore, examining the heterogeneous patterns of outpatient referral steering among MA beneficiaries reveals increased steering towards more costly "home visits". This type of care is particularly relevant in the context of differential coding, where home visits are a common strategy for diagnosing more patients.

Then, leveraging the structure of our framework and its welfare implications, our estimates of differential coding result in a welfare loss due to an increased excess burden of public funds and an increase in diagnostic effort costs. Whereas welfare improvements in referrals can theoretically offset these impacts, our empirical evidence does not support this interpretation among MA beneficiaries. Furthermore, our estimates of the impact of pay-per-patient contract adoption among MA beneficiaries mirror our results for the acquisition in MA. This is consistent with our framework, which highlights how both acquisition and contract adoption can achieve the same incentive alignment between insurers and physicians.

Next, we shift our focus towards the impact of the acquisition on Commercial beneficiaries. Here, as expected, there is no change in diagnostic behavior, as there are no incentives for differential coding. However, our findings are consistent with model prediction 1, as the acquired practice steers referrals towards more cost-effective specialist care, yielding estimated savings of approximately $300 per inpatient referral, and $26 per outpatient referral. Then, as predicted by our framework, the acquisition leads to welfare gains from referral steering among Commercial beneficiaries.

Combining these implications across segments highlights that evaluating the overall impact of acquisition requires us to trade off the welfare losses in the MA segment with the gains in the Commercial segment. Furthermore, the magnitude of the welfare costs of differential coding in MA depends on the magnitude of the excess burden of public funds parameter κ. Conversely, the welfare gains from steering referrals to more cost-effective specialists depend on both the number of referrals provided and how we interpret the savings per referral. We control for quality and therefore interpret our estimates as capturing cost savings for the insurer. However, if differences in negotiated prices are the result of differential market power, these savings could be interpreted as transfers between insurers and specialists. In order to interpret these cost savings as efficiency gains, the variation in prices needs to reflect the variation in the resource cost of specialist care instead.

Lastly, although the impacts of the acquisition on patients' health are beyond the scope of this paper, they may be of considerable interest for patient welfare. Whereas insurers propose that integration allows for improvements from increased care coordination, critics are often concerned about the impact on beneficiaries of other insurers, who may be foreclosed. We do find some evidence of improvements in the appropriateness of ER care following the acquisition. However, these outcomes are fundamentally noisy. Furthermore, concerns of exclusion of other insurers' patients are mitigated by our finding that the acquired practice continues to care for these types of patients. Lastly, just as pay-per-patient contracts have been shown to generate incentives for more preventative care (Morenz et al., 2023), acquisition may achieve the same effect.

Broader Lessons. Although our empirical evidence is limited to the impacts of one acquisition and pay-per-patient contract adoption in one market, these phenomena are not unique to our setting and may therefore allow for some broader conclusions. Insurer acquisitions of primary care practices have become increasingly common over the past decade, increasing policymaker concern over their impacts (DOJ, 2024; FTC, 2024). The acquisition we study is strikingly similar to other acquisitions carried out in the same time period (Table F.1). Insurers appear to be targeting large yet independent practices, which tend to fall under the reporting limit (Wollman, 2019), making it challenging to assess the full extent of this consolidation. Furthermore, these practices often care for a substantial number of Medicare patients, making them attractive targets for insurance plans in the Medicare Advantage segment. The seemingly broad targeting of practices caring for Medicare Advantage patients, along with pay-per-contract adoption being exclusive to these patients, is consistent with the idea that the insurer's objective is to align incentives in the MA segment, leading to differential coding.

Our findings of differential coding at the acquired practice and at practices that adopt pay-per-patient contracts highlight the importance of jointly considering the regulation of acquisitions and contracting. Then, common policy levers such as blocking acquisitions may be ineffective at avoiding these adverse outcomes in the presence of contracting. Although the causal relationship between acquisitions and contracting is outside the scope of this paper, our findings, combined with existing survey evidence and theories of contracting, provide suggestive evidence for evaluating the selection into and interaction of acquisitions and contracting.

Pay-per-patient contracts are considered challenging to implement by independent practices. Indeed, in surveys, physicians report that the health IT and administrative costs of implementation are key precluders of adoption (S. O'Malley et al., 2024; Damberg et al., 2014). In addition, common-agency problems (Frandsen and Rebitzer, 2019) limit the ability of insurers to bear the cost of these investments for physicians and overcoming these challenges to adoption. Therefore, acquisition can be interpreted as facilitating contract adoption.[48] Indeed, in a survey, over half of physicians report the "ease of participation in risk-based payment models" as a key reason for acquisition (Kane, 2025). Through acquisition, the insurer may provide the acquired practice with the necessary technology and staff investments that physicians may not be able to undertake on their own. This further highlights how selection into acquisition and pay-per-patient contracts may differ.

Then, the effectiveness of blocking this type of acquisition depends not only on the impact of the acquisition itself, but also on which kinds of practices would adopt these contracts in the absence of acquisition. Furthermore, acquisition can allow for further efficiencies within the firm that do not arise from contracting. For instance, acquisition may reduce information asymmetries, allow for better care coordination, and lead to positive spillovers onto Commercial beneficiaries,

[48] The idea that frictions may preclude the adoption of contracts is not new. Indeed, the transaction cost theory of integration argues that integration can be a response to these frictions (Bresnahan and Levin, 2012)

such as the referral steering we find.

7 Conclusion

This paper investigates the consequences of vertical integration between insurers and primary care physicians. We analyze a 2017 insurer acquisition of a primary care practice in Colorado, focusing on its impact on physicians' behavior. Our empirical analysis reveals two distinct and opposing effects. In the Medicare Advantage segment, the integrated practice increased its diagnostic coding for the acquiring insurer's beneficiaries, increasing risk adjustment payments by up to $1,805 per patient per year. In contrast, in the Commercial segment, the practice steered patients to more cost-effective specialists, generating savings of approximately $300 per inpatient referral.

Taken together, these findings highlight a trade-off in this form of vertical integration. The differential coding in Medicare Advantage represents a transfer from taxpayers to the integrated firm with little evidence of corresponding improvements in patient care, suggesting a welfare loss. Conversely, referral steering in the Commercial market represents a potential efficiency gain, as the integrated firm is incentivized to reduce the resource costs of care. Evaluating the net impact of such acquisitions therefore requires weighing these competing outcomes across market segments.

Our findings further imply that integration is not the only path to incentive alignment. We find that non-acquired practices adopt pay-per-patient contracts with the insurer and then also engage in differential coding, generating revenue increases statistically indistinguishable from those at the acquired practice. This suggests that contracts can act as substitutes for ownership, as financial incentives can be effective through either ownership or contracts.

Our analysis has several limitations that offer avenues for future research. First, although our results suggest implications outside of our specific setting, we examine a single acquisition event. Then, the extent to which our findings generalize to other markets and insurers remains an open question. Second, while we provide suggestive evidence that increased diagnoses were not met with equivalent increases in treatment, a full analysis of the long-term impacts on patient health outcomes is beyond the scope of this paper and warrants further attention. Finally, our framework takes the choice of organizational form as given; modeling the firm's endogenous decision to integrate versus contract is a promising direction for future theoretical and empirical work. Furthermore, understanding the practice characteristics and motivations behind insurers' targeting of practices can provide valuable insights into their potential effects on differential coding.

In conclusion, our paper offers a nuanced perspective on insurer-physician integration. While antitrust authorities are right to be concerned, our results show that the effects are complex, varying significantly across market segments. Most importantly, our finding that contracting can replicate key behaviors implies that policies focused narrowly on blocking mergers may be insufficient. An effective regulatory framework must look beyond firm boundaries to address the underlying financial incentives that drive firm and physician behavior, whether those incentives are embedded in an employment contract or a physician reimbursement agreement.

References

Agha, Leila, Keith Marzilli Ericson, Kimberley H. Geissler, and James B. Rebitzer, "Team Formation and Performance: Evidence from Healthcare Referral Networks," *NBER Working Paper 24338*, 2018.

Alliance, Better Medicare, "Understanding Risk Adjustment in Medicare Advantage," White Paper, Better Medicare Alliance 2017.

Atalay, Enghin, Ali Hortaçsu, and Chad Syverson, "Vertical Integration and Input Flows," *American Economic Review*, April 2014, *104* (4), 1120–48.

Atul, La Forgia Ambar Gupta and Adam Sacarny, "Turbocharging Profits? Contract Gaming and Revenue Allocation in Healthcare," 2025.

Baker, Laurence C., M. Kate Bundorf, and Daniel P. Kessler, "The effect of hospital/physician integration on hospital choice," *Journal of Health Economics*, 2016, *50*, 1–8.

Berenson, Robert A. and Mary Jo Braid-Forbes, "Development and Structure of BETOS 2.0 with Illustrative Data," Technical Report, Urban Institute 2020.

Billings, John, Nina Parikh, and Tod Mijanovich, "Emergency Department Use: The New York Story," *Issue Brief, The Commonwealth Fund*, 2000.

Bresnahan, Timothy F and Jonathan D Levin, "Vertical Integration and Market Structure," Working Paper 17889, National Bureau of Economic Research March 2012.

Brigham, Michael Powell Frandsen and James B. Rebitzer, "Sticking Points: Common-agency Problems and Contracting in the US Healthcare System," *The RAND Journal of Economics*, 2019, *50* (2), 251–285.

Brot-Goldberg, Zarek C and Mathijs de Vaan, "Intermediation and vertical integration in the market for surgeons," *Working Paper*, 2018.

Brown, Jason, Mark Duggan, Ilyana Kuziemko, and William Woolston, "How Does Risk Selection Respond to Risk Adjustment? New Evidence from the Medicare Advantage Program," *American Economic Review*, October 2014, *104* (10), 3335–64.

Callaway, Brantly and Pedro H.C. Sant'Anna, "Difference-in-Differences with multiple time periods," *Journal of Econometrics*, 2021, *225* (2), 200–230.

Capps, Cory, David Dranove, and Christopher Ody, "The effect of hospital acquisitions of physician practices on prices and spending," *Journal of Health Economics*, 2018, *59*, 139–152.

Centers for Medicare and Medicaid Services, “2017 Medicare Advantage ratebook and Prescription Drug rate information.” Available at https://www.cms.gov/medicare/health-plans/medicareadvtgspecratestats/ratebooks-and-supporting-data-items/2017rates.

_, “National Plan and Provider Enumeration System (NPPES).” Historical files available from NBER at https://www.nber.org/research/data/national-plan-and-provider-enumeration-system-nppesnpi.

Chan, David C, Matthew Gentzkow, and Chuan Yu, “Selection with Variation in Diagnostic Skill: Evidence from Radiologists*,” *The Quarterly Journal of Economics*, January 2022, *137* (2), 729–783. _eprint: https://academic.oup.com/qje/article-pdf/137/2/729/43336043/qjab048.pdf.

Cho, Hyein, “Provider-Insurer Integration and Healthcare Delivery,” *Working Paper*, 2025.

Clemens, Jeffrey and Joshua D. Gottlieb, “Do Physicians’ Financial Incentives Affect Medical Treatment and Patient Health?,” *American Economic Review*, April 2014, *104* (4), 1320–49.

Cooper, Zack, Stuart V. Craig, Aristotelis Epanomeritakis, Matthew Grennan, Joseph R. Martinez, Fiona Scott Morton, and Ashley T. Swanson, “Are Hospital Acquisitions of Physician Practices Anticompetitive?,” *NBER Working Paper Series*, 2025, *Working Paper 34039.*

Crawford, Gregory S., Robin S. Lee, Michael D. Whinston, and Ali Yurukoglu, “The Welfare Effects of Vertical Integration in Multichannel Television Markets,” *Econometrica*, 2018, *86* (3), 891–954.

Cuesta, José Ignacio, Carlos E Noton, and Benjamin Vatter, “Vertical Integration and Plan Design in Healthcare Markets,” August 2024, (32833).

Curto, Vilsa, Liran Einav, Amy Finkelstein, Jonathan Levin, and Jay Bhattacharya, “Health Care Spending and Utilization in Public and Private Medicare,” *American Economic Journal: Applied Economics*, April 2019, *11* (2), 302–32.

Cutler, David M, Leemore Dafny, David Cm Grabowski, Steven Lee, and Christopher Ody, “Vertical Integration of Healthcare Providers Increases Self-Referrals and Can Reduce Downstream Competition: The Case of Hospital-Owned Skilled Nursing Facilities,” December 2020, (28305).

Cutler, David M., Mark B. McClellan, and Joseph P. Newhouse, *Medical Care Output and Productivity*, University of Chicago Press, 2001.

Dafny, Leemore, “Hospital Industry Consolidation — Still More to Come?,” *New England Journal of Medicine*, 2014, *370* (3), 198–199.

_ and David Dranove, “Regulatory Exploitation and Management Changes: Upcoding in the Hospital Industry,” *The Journal of Law and Economics*, 2009, *52* (2), 223–250.

Damberg, Cheryl L., "Health Care Consolidation: The Changing Landscape of the U.S. Health Care System," Testimony CTA2770-1, RAND Corporation 2023.

__ , Melony E. Sorbero, Susan L. Lovejoy, Grant Martsolf, Laura Raaen, and Daniel Mandel, "Measuring Success in Health Care Value-Based Purchasing," *RAND Corporation, Research Report*, 2014.

Decarolis, Francesco, "Medicare Part D: Are Insurers Gaming the Low Income Subsidy Design?," *American Economic Review*, April 2015, *105* (4), 1547–80.

__ , Maria Polyakova, and Stephen P. Ryan, "Subsidy Design in Privately Provided Social Insurance: Lessons from Medicare Part D," *Journal of Political Economy*, 2020, *128* (5), 1712–1752.

Diebel, Ayse Sera, "Vertical Integration in the U.S. Health Care Market: An Empirical Analysis of Hospital-Insurer Consolidation," *Working Paper*, 2018.

DOJ, "New Health Care Task Force Will Tackle Competition Problems in Health Care Markets," 2024. Available from https://shorturl.at/HeTOb.

Dranove, David, "Demand Inducement and the Physician/Patient Relationship," *Economic Inquiry*, 1988, *26* (2), 281–298.

Duflo, Esther, Michael Greenstone, Rohini Pande, and Nicholas Ryan, "Truth-telling by Third-party Auditors and the Response of Polluting Firms: Experimental Evidence from India*," *The Quarterly Journal of Economics*, September 2013, *128* (4), 1499–1545. _eprint: https://academic.oup.com/qje/article-pdf/128/4/1499/30635961/qjt024.pdf.

Einav, Liran, Amy Finkelstein, and Neale Mahoney, "Provider Incentives and Healthcare Costs: Evidence From Long-Term Care Hospitals," *Econometrica*, 2018, *86* (6), 2161–2219.

__ , __ , Raymond Kluender, and Paul Schrimpf, "Beyond Statistics: The Economic Content of Risk Scores," *American Economic Journal: Applied Economics*, April 2016, *8* (2), 195–224.

Eliason, Paul J, Benjamin Heebsh, Ryan C McDevitt, and James W Roberts, "How Acquisitions Affect Firm Behavior and Performance: Evidence from the Dialysis Industry*," *The Quarterly Journal of Economics*, November 2019, *135* (1), 221–267.

Fang, Hanming and Qing Gong, "Detecting Potential Overbilling in Medicare Reimbursement via Hours Worked," *American Economic Review*, February 2017, *107* (2), 562–91.

Frandsen, Brigham and James B. Rebitzer, "Structuring Incentives within Accountable Care Organizations," *The Journal of Law, Economics, and Organization*, May 2014, *31.*

FTC, "Docket No. ATR 102: Request for Information on Consolidation in Health Care Markets," Tech. Report, Federal Trade Commission 2024.

Gaynor, Martin, James B. Rebitzer, and Lowell J. Taylor, "Physician Incentives in Health Maintenance Organizations," *Journal of Political Economy*, 2004, *112* (4), 915–931.

Geruso, Michael and Timothy Layton, "Upcoding: Evidence from Medicare on Squishy Risk Adjustment," *Journal of Political Economy*, 2020, *128* (3), 984–1026.

Ghili, Soheil, Ben Handel, Igal Hendel, and Michael D Whinston, "Optimal Long-Term Health Insurance Contracts: Characterization, Computation, and Welfare Effects," *The Review of Economic Studies*, May 2023, *91* (2), 1085–1121.

Griffin, John M. and Dragon Yongjun Tang, "Did Credit Rating Agencies Make Unbiased Assumptions on CDOs?," *American Economic Review*, May 2011, *101* (3), 125–30.

Grossman, Sanford J. and Oliver D. Hart, "The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration," *Journal of Political Economy*, 1986, *94* (4), 691–719.

Gupta, Atul, "Impacts of Performance Pay for Hospitals: The Readmissions Reduction Program," *American Economic Review*, April 2021, *111* (4), 1241–83.

Handel, Ben and Kate Ho, "The industrial organization of health care markets," *Handbook of Industrial Organization*, 2021, *5* (1), 521–614.

Hart, Oliver and Jean Tirole, "Vertical Integration and Market Foreclosure," *Brookings Papers on Economic Activity*, 1990, *21* (1990 Microeconomics), 205–286.

Ho, Kate and Ariel Pakes, "Hospital Choices, Hospital Prices, and Financial Incentives to Physicians," *American Economic Review*, 2014, *104* (12), 3841–3884.

_ and _ , "Physician Payment Reform and Hospital Referrals," *American Economic Review: Papers Proceedings*, 2014, *104* (5), 200–205.

Hortaçsu, Ali and Chad Syverson, "Cementing Relationships: Vertical Integration, Foreclosure, Productivity, and Prices," *Journal of Political Economy*, 2007, *115* (2), 250–301.

Jacobson, Gretchen, Tricia Neuman, and Anthony Damico, "Medicare Advantage Plan Switching: Exception or Norm?," *Kaiser Family Foundation*, 2016.

Johnson, Garret, Zoe M. Lyon, and Austin Frakt, "Provider-Offered Medicare Advantage Plans: Recent Growth And Care Quality," *Health Affairs*, 2017, *36* (3), 539–547. PMID: 28264957.

Kane, Carol K., "Physician Practice Characteristics in 2024: Private Practices Account for Less Than Half of Physicians in Most Specialties," Report, American Medical Association 2025.

KFF, "Ten Things to Know About Consolidation in Health Care Provider Markets," April 2024. Available from https://shorturl.at/LsdJf.

Koch, Thomas G., Brett W. Wendling, and Nathan E. Wilson, "The Effects of Physician and Hospital Integration on Medicare Beneficiaries' Health Outcomes," *The Review of Economics and Statistics*, September 2021, *103* (4), 725–739. _eprint: https://direct.mit.edu/rest/article-pdf/103/4/725/1964959/rest_a_00924.pdf.

Liebman, Eli and Matthew T. Panhans, "Why do narrow network plans cost less?," *Health Economics*, 2021, *30* (10), 2437–2451.

Luco, Fernando and Guillermo Marshall, "The Competitive Impact of Vertical Integration by Multiproduct Firms," *American Economic Review*, July 2020, *110* (7), 2041–64.

Miller, Rupert G. Jr., "An Unbalanced Jackknife," *The Annals of Statistics*, 1974, *2* (5), 880–891.

Morenz, Anna M., Lingmei Zhou, Edwin S. Wong, and Joshua M. Liao, "Association Between Capitated Payments and Preventive Care Among U.S. Adults," *AJPM Focus*, 2023, *2* (3), 100116.

Newhouse, JP and AM Garber, "Geographic Variation in Health Care Spending in the United States: Insights From an Institute of Medicine Report," *JAMA*, 2013.

O'Malley, Ann S., Rumin Sarwar, Cindy Alvarez, and Eugene C. Rich, "Why Primary Care Practitioners Aren't Joining Value-Based Payment Models: Reasons and Potential Solutions," *Commonwealth Fund*, 2024.

PAI, "Updated Report: Hospital and Corporate Acquisition of Physician Practices and Physician Employment 2019-2023," April 2024. Available from `https://shorturl.at/jEpNf`.

Park, Sungchul, Brent A. Langellier, and David J. Meyers, "Differences between integrated and non-integrated plans in Medicare Advantage," *Health Services Research*, 2023, *58* (3), 560–568.

Perry, M. K., "Vertical Integration: Determinants and Effects," *Handbook of Industrial Organization*, 2008, *1.*

PitchBook Data Inc. `www.pitchbook.com`.

Rey, P. and J. Tirole, "A Primer on Foreclosure," *Handbook of Industrial Organization*, 2007, *3.*

Riordan, Michael H., "Competitive Effects of Vertical Integration," *Handbook of Antitrust Economics*, 2008.

Trish, Erin, Paul Ginsburg, Laura Gascue, and Geoffrey Joyce, "Physician Reimbursement in Medicare Advantage Compared With Traditional Medicare and Commercial Health Insurance," *JAMA internal medicine*, 2017, *177* (9), 1287–1295.

Williamson, Oliver E., "The Vertical Integration of Production: Market Failure Considerations," *The American Economic Review*, 1971, *61* (2), 112–123.

— , "The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting," *University of Illinois at Urbana-Champaign's Academy for Entrepreneurial Leadership Historical Research Reference in Entrepreneurship*, 1985.

Wollman, Thomas G., "Stealth Consolidation: Evidence from an Amendment to the Hart-Scott-Rodino Act," *American Economic Review: Insights*, 2019, *1* (1), 77–94.

WSJ, Wall Street Journal, Christopher Weaver, Anna Wilde Mathews, and Denny Jacob, "UnitedHealth Confirms Civil and Criminal Justice Department Probes," July 2025. Available from https://shorturl.at/SuBCU.

Appendices

A Setting and Data Appendix

Our primary goal is to evaluate the impacts of the acquisition on the acquired practice's care decisions for their patients. To do this, we need to identify the individual physicians who are a part of the acquired practice and the patients whose care they manage.

A.1 Identifying Physician Practices

We begin by grouping individual physicians into their respective practices, a process we apply to all physicians in the market. To define a practice, we group physicians who share the same "business mailing address". We use physicians' mailing address rather than their practice address as it better captures the parent organization when a single practice operates out of multiple locations.[49]

We implement this by converting all mailing addresses to geographic coordinates using the Google Maps API, which helps resolve typos and formatting inconsistencies. We then match physicians with coordinates that are identical up to the fourth decimal place (11-meter radius). To account for physicians changing practices over time, we leverage a historical record of the National Plan and Provider Enumeration System (NPPES) and perform a manual verification for the acquired practice.[50]

Out of the approximately 58,500 medical providers in our raw data, about 40,500 are listed as the billing provider in a claim. Our procedure matches 35,600 of them to 13,100 distinct groups over the 5 years in our sample. The average group consists of 3.12 physicians, but there is a large right tail, with only 31 groups listing more than 100 physicians, and the largest group in a given year consisting of 1,500 professionals. The acquired practice lies slightly over the 99^{th} percentile, at around 80 physicians per year. It is worth noting, however, that because not every medical professional in an organization shows up as the billing provider, these numbers are a lower bound on the number of physicians in a given practice. Furthermore, differing billing practices across groups, e.g., whether the organization shows up as the billing provider rather than an individual physician, make size comparison across organizations more complex.

A.2 Assigning Patients to Primary Care Practices

Next, we identify each patient's primary care practice (PCP). Our procedure follows the standard in the literature (Agha et al., 2018, e.g.,) We define a patient's assigned PCP as the practice where they received the most primary care visits each period: pre-acquisition (2015-2016) and post-acquisition (2017-2019). A primary care visit is defined as a claim where a service provider with a primary care taxonomy conducts a primary care procedure, as defined by the data provider.

[49] This approach follows the logic of the TIN-matching method of Baker et al. (2016), which also aims to identify the business entity.

[50] This measure is imperfect, as it relies on physicians regularly updating their information in the NPPES database.

We assign each patient to the practice with the highest visit count, using individual expenditure at each practice to break ties.

A.3 Constructing Patients' Risk Scores

One of our primary outcomes is the change in patient risk scores. We compute risk scores using the 2019 CMS-HCC software, which generates a score based on a patient's annual diagnoses and demographics (age and gender). This is the same software CMS uses to determine risk adjustment payments Medicare Advantage plans. While not designed for the Commercial population, it can still generate a risk score for these patients, which we use as a control.

The software converts individual diagnoses into indicators for Hierarchical Condition Categories (HCC), which are aggregated into the risk score via a weighted sum.[51] The higher the number of HCCs, which we henceforth refer to as "effective diagnoses", the higher a patient's risk score is expected to be, controlling for age and gender (Appendix Figure **??**). However, because different "effective diagnoses" have different weights, two individuals with the same number of effective diagnoses may differ in their risk score.

To isolate the diagnostic behavior of the acquired practice, we construct two versions of the risk score. First, we compile a comprehensive list of all diagnoses for each patient each year, cross-walking ICD-9 to ICD-10 codes where necessary.[52] Then, we use all of a patient's diagnoses from any provider to construct R^{ALL}. Lastly, we restrict to only the diagnoses recorded on claims submitted by physicians at the acquired practice to compute $R^{ACQ\ PCP}$. This isolates the portion of the risk score directly attributable to the acquired practice's diagnostic decisions.

B Conceptual Framework Appendix

Equilibrium Choices

Baseline.

We start by solving for referrals and diagnoses in the absence of the acquisition, i.e., our baseline case. First, we restate our equilibrium referrals, before deriving the expression for the diagnostic cutoffs.

Recall the optimal referral choice under no effort but with diagnosis, as required for a referral:

$$\frac{v_\theta^{1\prime}(r_\theta^*(e_\theta = 0))}{v_\theta^{0\prime}(r_\theta^*(e_\theta = 0))} = -\frac{p_\theta}{1 - p_\theta}$$

Furthermore, under $e_\theta = 1$, the optimal referral satisfy $v_\theta^{1\prime}(r_\theta^*(e_\theta = 1)) = 0$.

Next, we solve for the PCP's diagnostic choices, taking their referral choices as given by the above first-order conditions. The PCP chooses whether to exert effort and to diagnose their pa-

[51]The HCC weights are normalized such that the average FFS enrollee has a risk score of 1. The risk score formula further includes interactions across HCCs.

[52]The software takes in ICD-10 codes. Because our sample includes the transition from ICD-9 to ICD-10, we first need to crosswalk all the ICD-9 codes to their corresponding ICD-10.

tient, with observed probability of being healthy p_θ, to maximize their utility in Equation 1. Their optimal decisions are characterized by thresholds of observed probability of their patient's health. We further assume that the PCP's utility under each effort and diagnostic choice are monotonic in their patients observed health signal. We derive the following decision rule:

$$\{e_\theta^*, d_\theta^*\} = \begin{cases} \{0,1\} & \text{if } p_\theta < \frac{v_\theta^1(r_\theta^*(1)) - v_\theta^1(r_\theta^*(0)) - (\omega - \gamma_\theta)}{\alpha_0 + \beta_\theta + v_\theta^1(r_\theta^*(1)) - v_\theta^1(r_\theta^*(0)) + v_\theta^0(r_\theta^*(0))} \equiv \underline{p}_\theta^* \\ \{1, \cdot\} & \text{if } \underline{p}_\theta^* \leq p_\theta \leq 1 - \frac{\omega - \gamma_\theta}{-\alpha_1 + \beta_\theta + v_\theta^1(r_\theta^*(1))} \equiv \bar{p}_\theta^* \\ \{0,0\} & \text{if } p_\theta > \bar{p}_\theta^* \end{cases} \tag{B.12}$$

After Acquisition.

As before, we start by solving for optimal referrals and compute the first-order condition under acquisition. We start by solving under $e_\theta = 0$ and $d_\theta = 1$ as follows:

$$\max_{r_\theta} \; v_\theta^0(r_\theta)p_\theta + v_\theta^1(r_\theta)(1 - p_\theta) - \lambda s_\theta(r_\theta)$$

We then compute the first-order condition for referrals as follows, where $r_\theta^{\mathcal{I}}$ represents the optimal referral choice:

$$v_\theta^{1\prime}(r_\theta^{\mathcal{I}}(0))(1 - p) + v_\theta^{0\prime}(r_\theta^{\mathcal{I}}(0))p = \lambda s_\theta'(r_\theta^{\mathcal{I}}(0)).$$

Then, we solve for the optimal referral intensity under $e_\theta = 1$. In this case, the PCP perfectly observes the patient's health status and may only refer a diagnosed patient. Their payoff from referral is then $v_\theta^1(r_\theta) - \lambda s_\theta(r_\theta)$, and the optimal referrals satisfy $v_\theta^{1\prime}(r_\theta^{\mathcal{I}}(1)) = \lambda s_\theta'(r_\theta^{\mathcal{I}}(1))$. Note that, as before, the optimal choice of referrals is higher when the PCP exerts diagnostic effort, that is, $r_\theta^{\mathcal{I}}(1) > r_\theta^{\mathcal{I}}(0)$.

Then, we solve for the PCP's diagnostic choices as before, taking their referral choices as given by the above first-order conditions. This yields the following decision rule under acquisition:

$$\{e_\theta^{\mathcal{I}}, d_\theta^{\mathcal{I}}\} = \begin{cases} \{0,1\} & \text{if } p_\theta < \frac{v_\theta^1(r_\theta^{\mathcal{I}}(1)) - v_\theta^1(r_\theta^{\mathcal{I}}(0)) + \lambda(s_\theta(r_\theta^{\mathcal{I}}(0)) - s_\theta(r_\theta^{\mathcal{I}}(1))) - (\omega - \gamma_\theta + \lambda\gamma_\theta)}{\alpha_0 + \beta_\theta + v_\theta^1(r_\theta^{\mathcal{I}}(1)) - v_\theta^1(r_\theta^{\mathcal{I}}(0)) + v_\theta^0(r_\theta^{\mathcal{I}}(0)) + \lambda(p_{MA}^R - \beta_\theta - s_\theta(r_\theta^{\mathcal{I}}(1)))} \equiv \underline{p}_\theta^{\mathcal{I}} \\ \{1, \cdot\} & \text{if } \underline{p}_\theta^{\mathcal{I}} \leq p_\theta \leq 1 - \frac{\omega - \gamma_\theta + \lambda\gamma_\theta}{-\alpha_1 + \beta_\theta + v_\theta^1(r_\theta^{\mathcal{I}}(1)) + \lambda(p_{MA}^R - \beta_\theta - s_\theta(r_\theta^{\mathcal{I}}(1)))} \equiv \bar{p}_\theta^{\mathcal{I}} \\ \{0,0\} & \text{if } p_\theta > \bar{p}_\theta^{\mathcal{I}} \end{cases} \tag{B.13}$$

Comparative Statics.

Referral Impact.

Comparing the optimal referral choice before and after the acquisition reveals a decrease in referrals both in the presence and absence of PCP effort. Furthermore, this decrease in referrals is larger as incentives become more aligned. Lastly, if integration moves the PCP from providing effort into not providing effort, referrals are still decreasing. However, if the opposite is true, that is, under integration the PCP starts to exert effort, referrals may either increase or decrease depending on the parametric assumptions. Namely, for referrals to still be decreasing in the acquisition, it must be that $\lambda s'_\theta(r_\theta) > \frac{-p}{1-p} v_\theta^{0\prime}(r_\theta)$. Intuitively, this means that the marginal cost of referrals for the PCP is greater than the marginal benefit for healthy patients, weighed by their relative probability.

Diagnostic Impact.

To evaluate how the incentive alignment from the acquisition impacts the PCP's choices of effort and diagnoses, we compare the cutoffs before and after the acquisition. For the purposes of this discussion, we focus on the MA segment. Our model predicts that the PCP becomes more likely to provide a diagnosis in the absence of effort, however, the impact on effort is ambiguous. Regardless, under our assumptions, our model predicts an increase in the probability of diagnosis, which we refer to as *differential coding*. Intuitively, this is caused by the acquisition increasing the value of providing diagnoses for the PCP. This leverages the assumption that the insurers' payoffs are increasing with diagnoses, i.e. $p^R \mathbb{1}_{\{\theta=MA\}} - \beta_\theta - s_\theta(r_\theta^{\mathcal{I}}) > 0$.

We start by analyzing how incentive alignment yields an increase in the probability of providing a diagnosis in the absence of effort. This is equivalent to analyzing the impact on the lower bound, $\underline{p}_\theta^{\mathcal{I}}$, for patients who are more likely to be unhealthy. Under some parametric assumptions, outlined below, we find :

$$\frac{\partial \underline{p}_\theta^{\mathcal{I}}}{\partial \lambda} = \frac{\overbrace{(s(r_\theta^{\mathcal{I}}(0)) - s(r_\theta^{\mathcal{I}}(1)) - \gamma)}^{<0} \overbrace{[\text{den}]}^{<0} - \overbrace{(\omega - \gamma_\theta + \lambda\gamma_\theta)}^{>0} \overbrace{[\text{num}]}^{<0}}{[\text{den}]^2} > 0$$

Intuitively, as the cost of effort and the value of diagnoses both increase, the PCP is more likely to directly diagnose their patients and not exert effort among patients who are more likely to be unhealthy. This impact both increases the probability of diagnosis and decreases the probability of effort.

Then, we focus on the impact of incentive alignment on the upper bound , $\bar{p}_\theta^{\mathcal{I}}$, for patients whose probability of being healthy is on the higher end. Whereas an increase in the cost of effort pushes the upper bound down, the increased value of diagnosing patients pushes it up. Then, whether the PCP is more or less likely to exert effort for healthier seeming patients depends on this trade-off. To evaluate this analytically, we compute the following derivative:

$$\begin{aligned}\frac{\partial \bar{p}_\theta^{\mathcal{I}}}{\partial \lambda} &= -\frac{(\overbrace{\gamma_\theta}^{>0}) \overbrace{[\text{den}]}^{>0} - \overbrace{(p_{MA}^R - \beta_\theta - s(r_\theta^{\mathcal{I}}(1)))}^{>0} \overbrace{[\text{num}]}^{>0}}{[\text{den}]^2} \\ &\propto (\omega - \gamma)(p_{MA}^R - \beta_\theta - s(r_\theta^{\mathcal{I}}(1))) - \gamma(-\alpha_1 + \beta_\theta + v_\theta^1(r_\theta(1)))\end{aligned}$$

Then, if $\frac{\omega}{\gamma} > \frac{p_{MA}^R - \alpha_1 + v_\theta^1(r_\theta(1)) - s(r_\theta^{\mathcal{I}}(1))}{p_{MA}^R - \beta_\theta - s(r_\theta^{\mathcal{I}}(1))}$, the upper bound is increasing in incentive alignment. This can be interpreted as the upper bound increasing when the cost of effort relative to the payment for effort is larger than the increased cost of referrals from having more people be eligible to be referred from increased diagnosis relative to the value of providing the extra diagnosis.

Lastly, combining these predictions, we can evaluate whether effort is predicted to increase or decrease. We define the probability of effort as the probability that the patient's health signal p_θ is in the interval $[\underline{p}_\theta, \bar{p}_\theta]$. Then, as both the lower and upper bound are predicted to increase, the probability of exerting effort is ambiguous and depends on which effect dominates.

Welfare Impact.

We leverage our expression for welfare in Equation 2 to outline the different impacts of the acquisition by computing the change in welfare before and after the acquisition.

$$
\begin{aligned}
\Delta W_\theta = & \underbrace{\alpha_0 p_\theta \left(\underline{p}_\theta^{\mathcal{I}} - \underline{p}_\theta^{*}\right) - \alpha_1 (1-p_\theta)\left(\bar{p}_\theta^{\mathcal{I}} - \bar{p}_\theta^{*}\right)}_{i.\Delta \text{Diagnostic precision}} - \underbrace{\kappa p^R \mathbb{1}\{\theta = MA\}\left(p_\theta\left(\underline{p}_\theta^{\mathcal{I}} - \underline{p}_\theta^{*}\right) + (1-p_\theta)\left(\bar{p}_\theta^{\mathcal{I}} - \bar{p}_\theta^{*}\right)\right)}_{ii.\Delta \text{Excess burden of public funds}} \\
& - \underbrace{\omega\left(-\underline{p}_\theta^{\mathcal{I}} + \underline{p}_\theta^{*} + \bar{p}_\theta^{\mathcal{I}} - \bar{p}_\theta^{*}\right)}_{iii.\Delta \text{Effort cost}} + \underbrace{\underline{p}_\theta^{\mathcal{I}}\left(\mathbb{E}_{p_\theta}\{v_\theta^i(r_\theta^{\mathcal{I}}(0)\} - c_\theta(r_\theta^{\mathcal{I}}(0))\right) - \underline{p}_\theta^{*}\left(\mathbb{E}_{p_\theta}\{v_\theta^i(r_\theta^{*}(0)\} - c_\theta(r_\theta^{*}(0))\right)}_{iv.\Delta \text{Average referral utility}} \\
& \underbrace{(1-p_\theta)\left(\left(\bar{p}_\theta^{\mathcal{I}} - \underline{p}_\theta^{\mathcal{I}}\right)\left(v_\theta^1(r_\theta^{\mathcal{I}}(1) - c_\theta(r_\theta^{\mathcal{I}}(1))\right) - \left(\bar{p}_\theta^{*} - \underline{p}_\theta^{*}\right)\left(v_\theta^1(r_\theta^{*}(1) - c_\theta(r_\theta^{*}(1))\right)\right)}_{v.\Delta \text{Referral utility among unhealthy}} \qquad \text{(B.14)}
\end{aligned}
$$

We begin by examining the welfare implications of increased diagnoses in the first three terms. Term *i* captures the change in welfare from changes in the precision of diagnostic behavior, although we note that we cannot empirically speak to this impact as we don't observe patients' true underlying health status. On the one hand, the increase in diagnoses among healthy patients generates a welfare loss. On the other hand, the extra diagnoses among sick patients who appeared healthier and weren't being tested before increase welfare. The overall impact of this depends on how we trade off between over- and under-diagnosing, that is, the relative magnitudes of α_0 and α_1, as well as the relative sizes of each of these populations.[53] Term *ii* captures the loss in welfare from the excess burden of government expenditures as the PCP provides more diagnoses. Lastly, term *iii* captures the change in welfare from the change in diagnostic effort. Whether this increases or decreases welfare depends on the total effect on effort, which is ambiguous in our model.

Then, we examine the impact on welfare from changes in referral behavior in terms *iv* and *v*. Both of these terms trade off the extensive margin change of referring more individuals as they are diagnosed and become eligible, and the intensive margin of reducing referrals after integration. Insofar as the reduction in referrals generates larger cost savings relative to the decreased patient utility, then the intensive margin reduction in referrals is beneficial for welfare. Furthermore, note that the scope for this to be beneficial depends on the variation in specialist cost. Lastly, the

[53] This depends on the magnitude of the increase in the cutoff at each end, and, under different distributions, this also depends on the mass of people near each cutoff.

increase in the extensive margin may be beneficial or detrimental for welfare for different groups of patients. Among healthy patients, this decreases welfare, whereas it improves it for unhealthy patients.

Pay-per-Patient Contracts. So far, we have defined payments to the PCP broadly, such that the PCP may receive payments β_θ for their diagnostic provision and payments γ_θ for their effort provision. By defining these parameters in different ways, our framework can reflect either of two payment arrangements: fee-for-service and pay-per-patient contracts. Furthermore, moving from fee-for-service towards pay-per-patient contracts can be seen as generating the same incentive changes as integration.

Under fee-for-service, PCPs are generally not paid per diagnosis, but rather per test or procedure. In contrast, under pay-per-patient contracts, PCPs are generally paid a fixed amount per patient, as a function of their diagnoses. Furthermore, pay-per-patient contracts often include a cost-sharing feature, whereby PCPs bear a share of the cost of specialist care. We can write the change in the PCP's objective function when going from fee-for-service to pay-per-patient as follows, where we define the baseline PCP utility as that under fee-for-service:

$$u^{PPP}(\cdot) = \begin{cases} u^{PCP}(0,0,\cdot) & \text{if } e_\theta = 0 \; \& \; d_\theta = 0 \\ u^{PCP}(0,1,r_\theta) + \left(\beta_\theta^{PPP} - \beta_\theta - \zeta^{PPP} s_\theta(r_\theta)\right) & \text{if } e_\theta = 0 \; \& \; d_\theta = 1 \\ u^{PCP}(1,\cdot,r_\theta) + (1-p_\theta)(\beta_\theta^{PPP} - \beta_\theta - \zeta^{PPP} s_\theta(r_\theta)) - (\gamma_\theta^{PPP} - \gamma_\theta) & \text{if } e_\theta = 1 \end{cases}$$

Then, a pay-per-patient contract can replicate the incentives from acquisition by designing the payments to match the change from an independent to an integrated PCP.

Including Treatment Provision. We can extend this model to allow for the PCP to also choose treatment provision $\tau = \{0,1\}$, with effort cost of treatment ω_τ, payment γ_τ, and value of treatment $\mu_0 < \mu_1$ for healthy and unhealthy patients respectively. We omit the θ notation for simplicity and rewrite the PCP's utility after acquisition as follows:

$$v^{ACQ\ PCP}(\cdot|\lambda) = \begin{cases} u^{ACQ\ PCP}(0,0,\cdot) & \text{if } e = 0 \; \& \; d = 0 \; \& \; \tau = 0 \\ u^{ACQ\ PCP}(0,0,\cdot) + p\mu_0 + (1-p)\mu_1 - (\omega_\tau - \gamma_\tau + \lambda\gamma_\tau) & \text{if } e = 0 \; \& \; d = 0 \; \& \; \tau = 1 \\ u^{ACQ\ PCP}(0,1,r) & \text{if } e = 0 \; \& \; d = 1 \; \& \; \tau = 0 \\ u^{ACQ\ PCP}(0,1,r) + p\mu_0 + (1-p)\mu_1 - (\omega_\tau - \gamma_\tau + \lambda\gamma_\tau) & \text{if } e = 0 \; \& \; d = 1 \; \& \; \tau = 1 \\ u^{ACQ\ PCP}(1,\cdot,r) & \text{if } e = 1 \\ u^{ACQ\ PCP}(1,\cdot,r) + (1-p)(\mu_1 - (\omega_\tau - \gamma_\tau + \lambda\gamma_\tau)) & \text{if } e = 1 \; \& \; \tau = 1 \end{cases}$$

We impose a similar monotonicity assumption on treatment, which implies that the PCP does not treat undiagnosed patients. Furthermore, when the PCP exerts effort and is certain of their patient's condition, they always diagnose and treat unhealthy patients. Then, the modified problem is characterized by three cutoffs. For patients who appear to be unhealthy enough, the PCP diagnoses and treats them while not exerting diagnostic effort. For patients who appear to be a little healthier, the PCP diagnoses them but does not treat them nor exerts effort. Lastly, for patients where there is more uncertainty, the PCP will exert diagnostic effort and diagnose and treat them appropriately.

We characterize the cutoffs as follows:

$$\{e^{\mathcal{I}}, d^{\mathcal{I}}, \tau^{\mathcal{I}}\} = \begin{cases} \{0,1,1\} & \text{if } p < \frac{\mu_1 - (\omega_\tau - \gamma_\tau + \lambda\gamma_\tau)}{\mu_1 - \mu_0} \equiv \underline{\underline{p}}^{\mathcal{I}} \\ \{0,1,0\} & \text{if } \underline{\underline{p}}^{\mathcal{I}} \leq p < \frac{v^1(r^{\mathcal{I}}(1)) - v^1(r^{\mathcal{I}}(0)) + \lambda(s(r^{\mathcal{I}}(0)) - s(r^{\mathcal{I}}(1))) - (\omega - \gamma + \lambda\gamma) + \mu_1 - (\omega_\tau - \gamma_\tau + \lambda\gamma_\tau)}{\alpha_0 + \beta + v^1(r^{\mathcal{I}}(1)) - v^1(r^{\mathcal{I}}(0)) + v^0(r^{\mathcal{I}}(0)) + \lambda(p^R_{MA} - \beta - s(r^{\mathcal{I}}(1)))} \equiv \underline{p}^{\mathcal{I}} \\ \{1,\cdot,\cdot\} & \text{if } \underline{p}^{\mathcal{I}} \leq p \leq 1 - \frac{\omega - \gamma + \lambda\gamma}{-\alpha_1 + \beta + v^1(r^{\mathcal{I}}(1)) + \lambda(p^R_{MA} - \beta - s(r^{\mathcal{I}}(1))) + \mu_1 - (\omega_\tau - \gamma_\tau + \lambda\gamma_\tau)} \equiv \bar{p}^{\mathcal{I}} \\ \{0,0,0\} & \text{if } p > \bar{p}^{\mathcal{I}} \end{cases}$$

To evaluate how they change with the acquisition, we compute their derivatives with respect to λ as before. We start by evaluating how the probability of not exerting effort but diagnosing and treating patients changes with the acquisition:

$$\frac{\partial \underline{\underline{p}}^{\mathcal{I}}}{\partial \lambda} = \frac{-\gamma_\tau}{\mu_1 - \mu_0} < 0$$

Then, as the PCP's incentives become aligned with the PCP, they become less likely to treat patients who appear sicker. Next, we evaluate how the probability of not exerting effort nor treating while diagnosing patients changes with the acquisition:

$$\frac{\partial \underline{p}^{\mathcal{I}}}{\partial \lambda} = \frac{\overbrace{(s(r^{\mathcal{I}}(0)) - s(r^{\mathcal{I}}(1)) - \gamma - \gamma_\tau)}^{<0} \overbrace{[\text{den}]}^{<0} - (\overbrace{\omega - \gamma + \lambda\gamma}^{>0}) \overbrace{[\text{num}]}^{<0}}{[\text{den}]^2} > 0$$

As the acquisition aligns incentives, not only is the PCP less likely to exert diagnostic effort for these patients, but also less likely to provide them with treatment. Lastly, we evaluate how the probability of not exerting effort nor treating or diagnosing patients changes with the acquisition:

$$\begin{aligned} \frac{\partial \bar{p}^{\mathcal{I}}}{\partial \lambda} &= -\frac{(\overbrace{\gamma}^{>0}) \overbrace{[\text{den}]}^{>0} - (\overbrace{p^R_{MA} - \beta - s(r^{\mathcal{I}}(1)) - \gamma_\tau}^{>0}) \overbrace{[\text{num}]}^{>0}}{[\text{den}]^2} \\ &\propto (\omega - \gamma)(p^R_{MA} - \beta - s(r^{\mathcal{I}}(1)) - \gamma_\tau) - \gamma(-\alpha_1 + \beta + v^1(r(1)) + \mu_1 - \omega_\tau) \end{aligned}$$

Then, if $\frac{\omega}{\gamma} > \frac{p^R_{MA} - \alpha_1 + v^1(r(1)) - s(r^{\mathcal{I}}(1)) + \mu_1 - \omega_\tau}{p^R_{MA} - \beta - s(r^{\mathcal{I}}(1)) - \gamma_\tau}$, the upper bound is increasing in incentive alignment.

We compute the expected probability of treatment as follows:

$$\mathbb{E}_p(\tau) = \underline{\underline{p}} + (\bar{p} - \underline{p}) \left(1 - \frac{\bar{p} + \underline{\underline{p}}}{2}\right)$$

Recall the expression for the expected probability of diagnosis:

$$\mathbb{E}_p(d) = \underline{p} + (\bar{p} - \underline{p}) \left(1 - \frac{\bar{p} + \underline{p}}{2}\right)$$

Then, from $\underline{\underline{p}} \leq \underline{p}$, it must be that $\mathbb{E}(\tau) \leq \mathbb{E}(d)$, that is the PCP weakly diagnoses more than they provide treatment. Furthermore, computing the derivatives of these expressions with respect to λ yields the following expressions:

$$\begin{aligned}
\frac{\partial \mathbb{E}_p(d)}{\partial \lambda} &= \frac{\partial \bar{p}}{\partial \lambda}(1 - \bar{p}) + \frac{\partial \underline{p}}{\partial \lambda}\underline{p} > 0 \\
\frac{\partial \mathbb{E}_p(\tau)}{\partial \lambda} &= \underbrace{\frac{\partial \underline{\underline{p}}}{\partial \lambda}}_{<0} + \underbrace{\frac{\partial \bar{p}}{\partial \lambda}(1 - \bar{p})}_{>0} - \underbrace{\frac{\partial \underline{p}}{\partial \lambda}(1 - \underline{p})}_{>0}
\end{aligned}$$

Note that, although the sign of the second equation is undeterminate, even if treatment is increasing in incentive alignment, it increases by less than diagnoses do:

$$\frac{\partial \mathbb{E}_p(d)}{\partial \lambda} - \frac{\partial \mathbb{E}_p(\tau)}{\partial \lambda} = \underbrace{\frac{\partial \underline{p}}{\partial \lambda}}_{>0} - \underbrace{\frac{\partial \underline{\underline{p}}}{\partial \lambda}}_{<0} > 0$$

Assumptions. We assume that the PCP receives disutility from misdiagnosing their patient, $\alpha_0, \alpha_1 < 0$. Furthermore, the PCP receives non-negative payments for their diagnoses, $\beta_\theta \geq 0$, and payments to partially cover the cost of diagnostic effort, $\omega > \gamma_\theta \geq 0$. We further assume that the utility for referrals for unhealthy patients is greater than that of healthy patients, $v_\theta^1 > v_\theta^0$, and that it is increasing in referrals for unhealthy patients but decreasing for healthy patients but concave for all, $v_\theta^{1\prime} > 0, v_\theta^{0\prime} < 0$, and $v_\theta^{1\prime\prime}, v_\theta^{0\prime\prime} < 0$. We further impose that the cost of referrals for the insurer is increasing and convex in referrals $s'_\theta, s''_\theta > 0$.

We impose that the insurer would prefer to have more diagnoses among their MA patients: $p^R - \beta_\theta - s_\theta(r_\theta(1)) > p^R - \beta_\theta - s_\theta(r_\theta(0)) > 0$. We also assume that the penalty for misdiagnosing healthy patients (α_0) is large enough, that is, $\alpha_0 < -[\beta + v_\theta^1(r_\theta(1)) - v_\theta^1(r_\theta(0)) + v_\theta^0(r_\theta(0)) + \lambda(p^R - \beta_\theta - s_\theta(r_\theta(1)))]$, and that the gain from increased referrals among healthy patients when exerting effort is smaller than the cost of effort $v_\theta^1(r_\theta(1)) - v_\theta^1(r_\theta(0)) < \omega - \gamma_\theta$.

Relaxing Assumptions. If we don't assume that the PCP's choices are monotonically ordered, this yields an extra boundary to consider, between providing a diagnosis or not in the absence of effort. Under no effort, the PCP provides a diagnosis if the expected payoff of diagnosis is larger than that of no diagnosis. Then, the PCP diagnoses their patient only if the patient's observed probability of being healthy meets the following condition:

$$p_\theta \leq \frac{\alpha_1 - \beta_\theta - v_\theta^1(r_\theta^*(0))}{\alpha_0 + \alpha_1 + v_\theta^0(r_\theta^*(0)) - v_\theta^1(r_\theta^*(0))} \equiv \tilde{p}_\theta^*$$

Under the acquisition, this threshold is modified as follows:

$$\begin{aligned} p_\theta \leq & \frac{\alpha_1 - \beta_\theta - v_\theta^1(r_\theta^{\mathcal{I}}(0)) - \lambda(p_{MA}^R - \beta_\theta - s_\theta(r_\theta^{\mathcal{I}}(0)))}{\alpha_0 + \alpha_1 + v_\theta^0(r_\theta^{\mathcal{I}}(0)) - v_\theta^1(r_\theta^{\mathcal{I}}(0))} \\ \leq & \tilde{p}_\theta^* - \frac{\lambda(p_{MA}^R - \beta_\theta - s_\theta(r_\theta^{\mathcal{I}}(0)))}{\alpha_0 + \alpha_1 + v_\theta^0(r_\theta^{\mathcal{I}}(0)) - v_\theta^1(r_\theta^{\mathcal{I}}(0))} \equiv \tilde{p}_\theta^{\mathcal{I}}, \end{aligned}$$

where it is straightforward to see that $\tilde{p}_\theta^{\mathcal{I}} \geq \tilde{p}_\theta^*$, that is, in the absence of effort, the PCP becomes more likely to provide a diagnosis.

In turn, this makes the change in the probability of diagnosis more complex:

$$\frac{\partial \mathbb{E}_{p_\theta}(d_\theta)}{\partial \lambda} = \begin{cases} \frac{\partial \tilde{p}_\theta}{\partial \lambda} + \frac{\partial \bar{p}_\theta}{\partial \lambda}(1 - \bar{p}_\theta) - \frac{\partial \underline{p}_\theta}{\partial \lambda}(1 - \underline{p}_\theta) \\ \text{if } \tilde{p}_\theta < \underline{p}_\theta \\ \frac{\partial \bar{p}_\theta}{\partial \lambda}(1 - \bar{p}_\theta) + \frac{\partial \underline{p}_\theta}{\partial \lambda}\underline{p}_\theta \\ \text{if } \underline{p}_\theta \leq \tilde{p}_\theta \leq \bar{p}_\theta \\ \frac{\partial \tilde{p}_\theta}{\partial \lambda} - \frac{\partial \bar{p}_\theta}{\partial \lambda}\bar{p}_\theta + \frac{\partial \underline{p}_\theta}{\partial \lambda}\underline{p}_\theta \\ \text{if } \tilde{p}_\theta \geq \bar{p}_\theta \end{cases}$$

Note that, in our main analysis, we consider the case where $\tilde{p}_\theta \in [\underline{p}_\theta, \bar{p}_\theta]$.

C Differential Coding

C.1 Margins of Differential Coding in Risk Space

The underlying thought experiment is moving the extra patient with any diagnosis or the extra diagnosis to have the impact of the *average* diagnosis. In so far as we think that the marginal individual being diagnosed with anything (or the marginal diagnosis given to a patient) are different from the average, this is not capturing that difference. It is nevertheless still a useful way to compare the magnitude of the impacts to each other. We do this for the two outcomes that are not in risk space: "Any Diagnosis" and the "Quantity" of diagnoses.

We start by outlining the implementation for "Any Diagnosis". For the treated group, we impute what the risk score of the average treated individual with any diagnoses vs. no diagnosis would be in the post periods. For the control group, use the realized average risk of individuals with any vs. no diagnoses. Mathematically, let $d = \mathbb{1}\{D_{it}^{ACQ\ PCP} > 0\} \in \{0, 1\}$ be an indicator for whether an individual has any diagnoses. The superscript C indicates the control group, the superscript T indicates the treatment group. We then redefine the outcome as follows:

$$\tilde{R}_t^C(D=d) = \frac{\sum_i \mathrm{R}_{it}^{ACQ\ PCP} \mathbb{I}_t\{i \in \text{Other Ins}, D=d\}}{\sum_i \mathbb{I}_t\{i \in \text{Other Ins}, D=d\}}$$

$$\tilde{R}_{2016}^T(D=d) = \frac{\sum_i \mathrm{R}_{i,2016}^{ACQ\ PCP} \mathbb{I}_{2016}\{i \in \text{Acq Ins}, D=d\}}{\sum_i \mathbb{I}_{2016}\{i \in \text{Acq Ins}, D=d\}}$$

$$\tilde{R}_t^T(D=d) = \tilde{R}_t^C(D=d) - \left(\tilde{R}_{2016}^C(D=d) - \tilde{R}_{2016}^T(D=d)\right)$$

Then, we implement this for the "Quantity" of diagnosis. For the treated group, we impute what the average risk per diagnosis would be and adjust for the likelihood of getting any diagnoses. Mathematically, we start by defining and imputing the average risk score per diagnosis as follows:

$$\widetilde{AR}_t^C = \frac{\sum_i \frac{R_{it}^{ACQ\ PCP}}{D_{it}^{ACQ\ PCP}} \mathbb{I}_t\{i \in \text{Other Ins}, D=1\}}{\sum_i \mathbb{I}_t\{i \in \text{Other Ins}, D=1\}}$$

$$\widetilde{AR}_{2016}^T = \frac{\sum_i \frac{R_{it}^{ACQ\ PCP}}{D_{i,2016}^{ACQ\ PCP}} \mathbb{I}_{2016}\{i \in \text{Acq Ins}, D=1\}}{\sum_i \mathbb{I}_{2016}\{i \in \text{Acq Ins}, D=1\}}$$

$$\widetilde{AR}_t^T = \widetilde{AR}_t^C - \left(\widetilde{AR}_{2016}^C - \widetilde{AR}_{2016}^T\right)$$

Then, define and impute the share of patients with any diagnosis as follows:

$$\tilde{S}_t^C = \frac{\sum_i \mathbb{I}_t\{i \in \text{Other Ins}, D=1\}}{\sum_i \mathbb{I}_t\{i \in \text{Other Ins}\}}$$

$$\tilde{S}_{2016}^T = \frac{\sum_i \mathbb{I}_{2016}\{i \in \text{Acq Ins}, D=1\}}{\sum_i \mathbb{I}_{2016}\{i \in \text{Acq Ins}\}}$$

$$\tilde{S}_t^T = \tilde{S}_t^C - \left(\tilde{S}_{2016}^C - \tilde{S}_{2016}^T\right)$$

Finally, we construct the outcome: $Y_{it} = \widetilde{AR}_t * \tilde{S}_t * D_{it}^{ACQ\ PCP}$.

D Referral Steering

D.1 Cost Index Construction

We define an observation as a claim for patient i, under insurance plan $g(i)$, at specialist j, in period t, for a specialist service group s. We divide specialist services into two categories, inpatient and outpatient services, and then into groups within each category. For inpatient claims, groups are defined at the diagnosis-related-group (DRG) claim classification. For outpatient claims, groups are defined by their BETOS code classification.[54] Then, within each group s, plan $g(i)$ and period t we retrieve the specialist cost index from the following linear regression:

[54] E.g., BETOS code M5C represents ophthalmology specialist evaluation and management claims.

$$\$ \text{ Amount}_{ijt} = \alpha_{g(i),s,t} + \text{Spe FE}_{g(i),j,s,t} + \mathbf{X}_{ijts} + \varepsilon_{ijts} \tag{D.15}$$

where $\mathbf{X}$ includes year, gender, pay-per-patient contract, and MA fixed effects, as well as controls for patient risk score.[55] Then, the specialist cost index varies at the specialist service, period, and insurance plan level. We run this regression with HC1 robust standard errors, demean the specialist fixed effects so that they are centered around 0, and then perform two corrections: a closed-form LOO correction and a hierarchical empirical Bayes shrinkage.

We follow Miller (1974) to perform a post-estimation, closed-form leave-one-out correction. Let $h_{ii} = [\mathbf{H}]_{ii} = \mathbf{x}_i^T \left(\mathbf{X}^T\mathbf{X}\right)^{-1}$ be the leverage score of observation i, computed as the diagonal of the ortho-projection matrix $\mathbf{H} = \mathbf{X}\left(\mathbf{X}^T\mathbf{X}\right)^{-1}\mathbf{X}^T$. Let $\hat{e}_i = y_i - \mathbf{x}_i^T\hat{\beta}$ be the regression residuals. The LOO estimated coefficient $\hat{\beta}^{(-i)}$ is then defined as:

$$\hat{\beta}^{(-i)} = \hat{\beta} - \frac{\left(\mathbf{X}^T\mathbf{X}\right)^{-1}\mathbf{x}_i\hat{e}_i}{1 - h_{ii}}$$

To reduce noise in our estimates, we shrink our estimates across two dimensions: within plan across specialists and within specialist across plans. This is a form of hierarchical shrinkage. Let $p \in [1, P]$ be a plan and $j \in [1, J]$ be a specialist (within a specialist service-period). We can visualize our specialist fixed matrix as follows, where a column is an insurance plan and a row is a specialist, as well as compute row and column means, $\bar{\beta}^j$ and $\bar{\beta}^p$ respectively, we denote each $\hat{\beta}_{pj}$'s standard errors as s_{pj}, and the matrix mean $\bar{\beta}_{pj}$:

$$\begin{array}{cc} \hat{\beta} = \begin{pmatrix} \hat{\beta}_{1,1} & \dots & \hat{\beta}_{1,J} \\ \hat{\beta}_{2,1} & \dots & \hat{\beta}_{2,J} \\ \vdots & \ddots & \vdots \\ \hat{\beta}_{P,1} & \dots & \hat{\beta}_{I,J} \end{pmatrix} & \begin{array}{l} \to \bar{\beta}_1^j \\ \\ \\ \\ \end{array} \\ \begin{array}{l} \qquad\quad\; \downarrow \\ \qquad\quad\; \bar{\beta}_1^p \end{array} & \end{array}$$

Note that this matrix is incomplete, as not all specialists interact with all plans. Then, let R be an indicator matrix for whether a plan-specialist cell in the $\hat{\beta}$ matrix is non-missing. Also note that this shrinkage assumes the missingness is random, which is likely not correct, but it lets us account for different sample sizes. We define the following quantities:

[55] Given that risk scores are endogenously increasing in our setting, we use pre-acquisition risk scores as our control

$$
\begin{aligned}
\tilde{s}_p^2 &\equiv \frac{\sum_j R_{pj} s_{pj}^2}{\sum_j R_{pj}} \\
\tilde{s}_j^2 &\equiv \frac{\sum_p R_{pj} s_{pj}^2}{\sum_p R_{pj}} \\
\tilde{s}_{pj}^2 &\equiv \frac{\sum_{pj} R_{pj} s_{pj}^2}{\sum_{pj} R_{pj}} \\
\hat{\xi}_{PJ} &= \frac{(PJ-1)\hat{V}(\bar{\beta}_{pj})}{(P-1)(J-1)} - \frac{J\hat{V}(\bar{\beta}_p)}{J-1} - \frac{P\hat{V}(\bar{\beta}_j)}{P-1} + \frac{\left[\frac{\sum_p \tilde{s}_p^2}{I(J-1)} + \frac{(J-1)\sum_p \tilde{s}_j^2}{J} - (PJ-1)\tilde{s}_{pj}^2\right]}{(P-1)(J-1)} \\
w_1 &= \left(\frac{s_{pj}^2}{\xi_{PJ} + s_{pj}^2}\right) \\
w_2 &= \left(\frac{\xi_{PJ}}{\xi_{PJ} + s_{pj}^2}\right)
\end{aligned}
$$

The final LOO Empirical Bayes Adjusted estimate can be written out as follows:

$$
\hat{\beta}_{pj}^{(-i),EB} = w_1\left(\bar{\beta}_p + \bar{\beta}_j\right) + w_2 \hat{\beta}_{pj}^{(-i)}
$$

E Appendix Figures

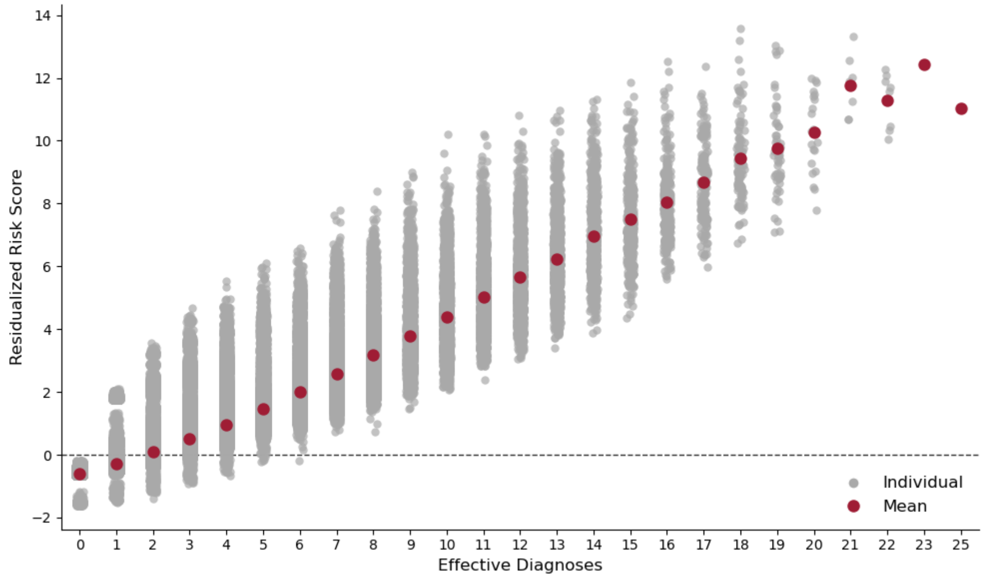

Figure E.1: Relationship between Risk Score and Effective Diagnoses

Notes: This figure illustrates the relationship between individuals' number of HCC codes, or "effective diagnoses", and their resulting risk score, controlling for age and gender. Risk scores are a weighted average of effective diagnoses. Differences in weights account for the variance within each number of diagnoses, as well as potential interactions between diagnoses. The average individual has 1.5 diagnoses.

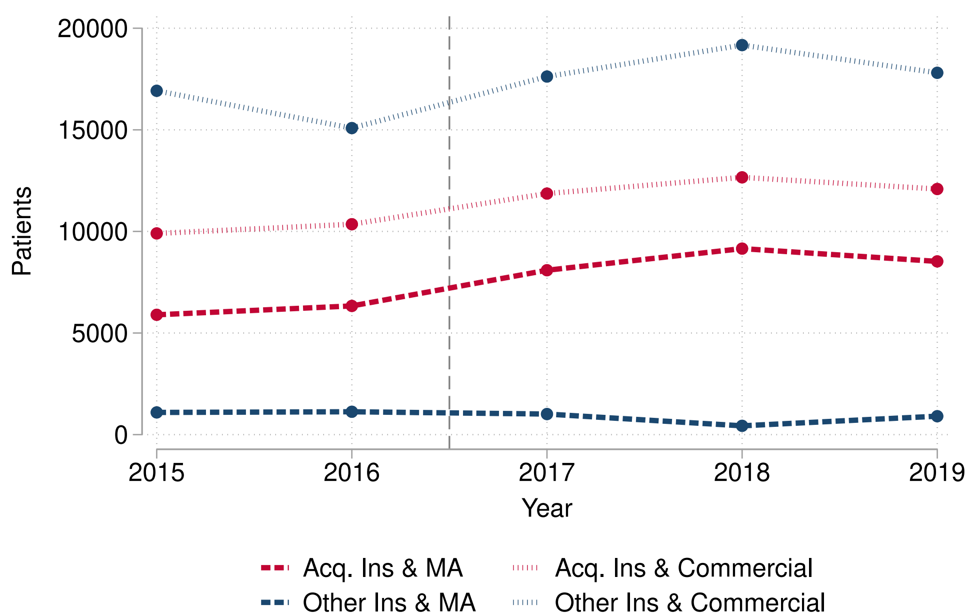

Figure E.2: Patient Composition at Acquired PCP

Notes: This figure plots the number of patients from each insurer and market segment at the acquired practice over time. The vertical line indicates the 2017 acquisition date.

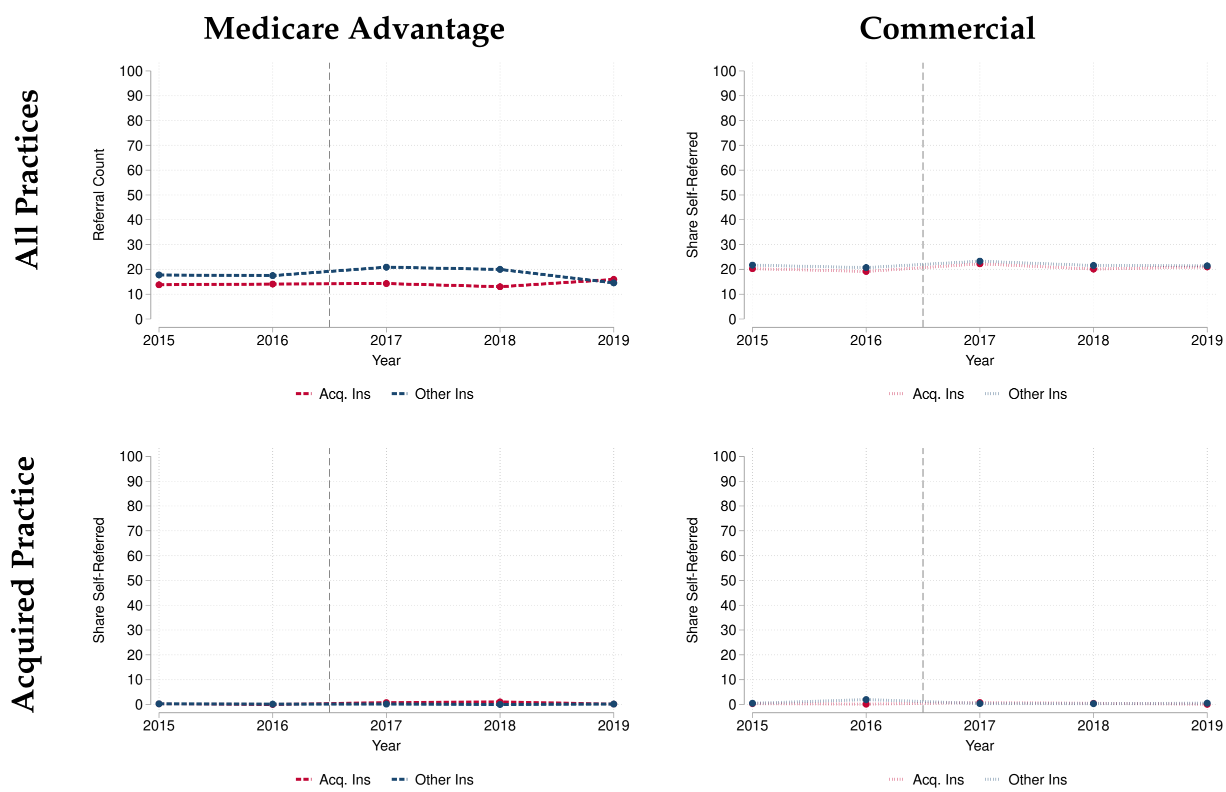

Figure E.3: Share of Within Practice Outpatient Referrals

Notes: The figures illustrate the share of outpatient referrals that are within a practice, that is, the amount of self steering. The figures on the first row show this across all practices, and the bottom row at the acquired practice only. The left column restricts to MA beneficiaries and the right column to Commercial beneficiaries. In red, we plot the share of self-steered referrals among beneficiaries of the acquiring insurer. In blue, we plot the share of self-steered referrals among beneficiaries of other insurers. The vertical line in each panel indicates the 2017 acquisition date

F Appendix Tables

Year	State	Patients	Providers	Locations	Care Model	Specialties	Transaction
2015	CT	350k	370	90	ACO	X	-
2016	NJ	200k	120	25	PCMH	X	-
2017	WA	320k	560	21	ACO	X	-
2017	CO	200k	100	20	ACO	X	-
2017	CA	20k	55	6	Managed Care	X	-
2018	MA	320k	500	25	ACO	X	$28M

Table F.1: Characteristics of Physician Acquisitions

This table lists six individual acquisitions carried out by the insurer in the mid to late 2010s. Acquisitions are identified using Pitchbook as well as press releases from the insurer/practices. Practice characteristics are self-reported from each practice's website prior to being acquired. For consistency across practices, we report the number of medical professionals, excluding administrative staff. The practices share several characteristics. First, these are all "private practices" before being acquired, i.e., they are not part of a hospital or health system. Second, they are large practices, employing a couple hundred medical professionals, and caring for 200,000 to 300,000 patients, including Medicare beneficiaries, across several practice locations. They are primary care practices but also provide some specialty care. Further, all of these practices adopt care models which share a focus on cost controls and care coordination, including Accountable Care Organizations (ACOs), managed care initiatives, and patient-centered medical homes (PCMHs). Lastly, we only observe the transaction value of the acquisition for one instance, at $28 million, which is below the reporting limit

	(1) $R^{ACQ\ PCP}$	(2) $R^{ACQ\ PCP}$	(3) $R^{ACQ\ PCP}$	(4) $R^{ACQ\ PCP}$
τ_{pre}	0.019	0.019	0.020	0.020
	(0.023)	(0.023)	(0.024)	(0.024)
τ_{post}	0.096	0.078	0.083	0.064
	(0.022)	(0.025)	(0.023)	(0.026)
Year FE	X	X	X	X
Observations	42,629	28,748	37,889	24,008
Patient Sample	All	Pre Acq	Post Acq	Balanced

Table F.2: Differential Coding w/o Individual FEs

Notes: The table presents the event study estimates for the change in patient risk scores across four populations. Across populations, we restrict to the sample to our Acquired PCP sample, that is, we keep only patients of the acquired practice. In column (1), we include all patients. In column (2), we include only those individuals who were patients of the acquired practice prior to acquisition. In column (3), we include only those individuals who were patients of the acquired practice after to acquisition. In column (2), we include only those individuals who were patients of the acquired practice both before and after the acquisition. An observation is a patient-year. Our estimates are computed using the event study design with time fixed effects from Callaway and Sant'Anna (2021). We do not include individual fixed effects. Control is not yet treated. τ_{post} captures the average treatment effect over the 3 years post acquisition.

In Patient Claims						
	Mean	SD	p25	p50	p75	N
Commercial	$1,726	$1,798	$640	$1,034	$2,211	8,562
MA	$944	$0.971	$234	$938	$1,249	2,915
Out Patient Claims						
	$Mean	SD	p25	p50	p75	N
Commercial	$50	$66	$9	$17	$61	148,641
MA	$30	$42	$10	$14	$27	106,980

Table F.3: Price Variation in Specialist Services

Notes: This table reports summary statistics for the total cost of inpatient and outpatient claims across market segments in 2016. Amounts are in dollars. Claims are restricted to those that feasibly result from a referral. Outpatient claims are restricted to BETOS Codes M4*, M5*, P1*-P5*, as these are the codes for which we conduct the referral exercise.

	Medicare Advantage								
	In Pat	Out Pat	Home Vis	Spe Off Vis	Major Proc	Cardio Proc	Ortho Proc	Eye Proc	Amb Proc
N	916	242,267	65,103	85,675	17,235	7,403	14,441	21,294	31,116
Mean	2.70e-2	3.64e+3	-7.83e-8	1.03e+4	-5.30e-7	3.92e-7	-8.76e-8	1.14e-6	-3.47e-8
SD	2.37e+0	4.74e+6	2.71e-7	7.98e+6	5.77e-6	2.76e-6	2.45e-6	1.64e-6	3.67e-6
p25	-1.26e-1	-2.25e-7	-1.67e-7	-3.40e-8	-4.01e-6	-4.84e-7	-1.30e-6	4.79e-7	-2.86e-6
p50	-1.65e-2	5.75e-8	-1.17e-7	1.68e-7	-4.28e-7	8.83e-7	-1.20e-7	1.34e-6	-3.66e-7
p75	6.98e-2	4.60e-7	-1.08e-8	3.59e-7	2.16e-6	1.80e-6	8.22e-7	1.99e-6	2.02e-6
	Commercial								
	In Pat	Out Pat	Home Vis	Spe Off Vis	Major Proc	Cardio Proc	Ortho Proc	Eye Proc	Amb Proc
N	1,387	174,310	6,631	63,494	28,487	4,022	15,176	10,983	45,517
Mean	8.11e-2	1.61e+0	4.23e+1	1.24e-7	8.92e-8	7.21e-7	1.77e-7	9.87e-7	1.53e-7
SD	3.52e+0	2.63e+4	1.35e+5	2.18e-6	8.20e-6	2.36e-6	2.49e-6	4.76e-6	6.30e-6
p25	-1.88e-1	-9.80e-7	-3.30e-7	-2.16e-7	-3.53e-6	2.45e-7	-1.19e-6	-2.92e-7	-3.23e-6
p50	0.00e+0	1.04e-7	-1.08e-7	1.05e-7	-2.76e-7	8.83e-7	2.47e-7	8.51e-7	8.50e-8
p75	1.92e-1	1.30e-6	-2.66e-8	5.29e-7	3.64e-6	2.54e-6	1.28e-6	1.66e-6	2.36e-6

Table F.4: Cost Index Summary Statistics

Notes: This table reports summary statistics for the cost index of inpatient and outpatient claims across market segments. Amounts are in dollars. Claims are restricted to those that feasibly result from a referral. Outpatient claims are restricted to BETOS Codes M4*, M5*, P1*-P5*, as these are the codes for which we conduct the referral exercise.

	Mean	SD	p25	p50	p75	N
Baseline						
Commercial	70	4,692	-253	0	231	1,387
MA	50	3,035	-149	-10	97	916
After Empirical Bayes Shrinkage						
	Mean	SD	p25	p50	p75	N
Commercial	81	3,517	-188	0	192	1,387
MA	27	2,366	-126	-17	70	916

Table F.5: Inpatient Cost Index Summary Pre and Post Shrinkage

Notes: This table reports summary statistics for the cost index of inpatient claims across market segments before and after shrinkage. Amounts are in thousands of dollars. Claims are restricted to those that feasibly result from a referral.

	Evaluation and Management	
	Home Visit	Specialist Visit
$\hat{\eta}_{kg(i)t}^{\text{Acq Ins \& Post}}$	-0.636	0.236
	(0.218)	(0.061)
$\hat{\eta}_{kg(i)t}^{\text{Acq Ins}}$	2.564	-0.089
	(0.175)	(0.029)
$\hat{\eta}_{kg(i)t}^{\text{Post}}$	0.524	0.080
	(0.060)	(0.010)
$\hat{\bar{\eta}}$	-2.014	-0.229
	(0.063)	(0.010)
$\hat{\beta}$	1.065	-0.061
	(0.065)	(0.013)
σ_{δ}	21.326	16.827
	(23.051)	(4.251)
Observations	754	10497
Δ Cost ($ per ref)	-1033	156
Referrals	289	1854

	Procedures				
	Major	Major Cardio	Major Ortho	Eye	Ambulatory
$\hat{\eta}_{kg(i)t}^{\text{Acq Ins \& Post}}$	0.246	0.334	-0.005	0.056	0.003
	(0.096)	(0.435)	(0.100)	(0.096)	(1.209)
$\hat{\eta}_{kg(i)t}^{\text{Acq Ins}}$	-0.102	-0.495	-0.101	-0.097	0.190
	(0.019)	(0.299)	(0.049)	(0.064)	(0.757)
$\hat{\eta}_{kg(i)t}^{\text{Post}}$	-0.169	-0.302	0.123	-0.025	-0.023
	(0.084)	(0.437)	(0.009)	(0.043)	(1.077)
$\hat{\bar{\eta}}$	-0.132	0.267	-0.073	0.161	-0.056
	(0.096)	(0.337)	(0.051)	(0.043)	(1.033)
$\hat{\beta}$	0.695	0.606	0.538	0.130	0.572
	(0.045)	(0.191)	(0.093)	(0.080)	(0.131)
σ_{δ}	15.912	16.588	13.965	17.539	17.176
	(21.724)	(20.774)	(7.491)	(9.050)	(285.552)
Observations	2924	404	1168	1484	4956
Δ Cost ($ per ref)	189	244	-5	75	3
Referrals	656	127	576	304	1150

Table F.6: Outpatient Referral Steering - Commercial - Random Effects

Notes: The table presents the model estimates for the change in the acquired PCP's cost sensitivity for inpatient referrals across outpatient specialist procedures for Commercial beneficiaries. An observation is a referral. We restrict our sample to our Acquired PCP sample, that is, we keep only patients of the acquired practice. The treated group are Commercial beneficiaries of the acquiring insurer at the acquired practice. The control group are Commercial beneficiaries of other insurers at the acquired practice. We include random effects in this specification.

	Evaluation and Management	
	Home Visit	Specialist Visit
$\hat{\eta}_{kg(i)t}^{\text{Acq Ins \& Post}}$	-0.757	0.117
	(0.123)	(0.049)
$\hat{\eta}_{kg(i)t}^{\text{Acq Ins}}$	1.105	0.015
	(0.093)	(0.031)
$\hat{\eta}_{kg(i)t}^{\text{Post}}$	0.209	-0.020
	(0.114)	(0.015)
$\hat{\eta}$	-0.862	0.054
	(0.061)	(0.011)
$\hat{\beta}$	0.362	-0.060
	(0.023)	(0.011)
Observations	754	10497
Δ Cost ($ per ref)	-2112	120
Referrals	289	1854

	Procedures				
	Major	Major Cardio	Major Ortho	Eye	Ambulatory
$\hat{\eta}_{kg(i)t}^{\text{Acq Ins \& Post}}$	0.211	-0.253	0.271	-0.058	0.064
	(0.091)	(0.116)	(0.046)	(0.035)	(0.056)
$\hat{\eta}_{kg(i)t}^{\text{Acq Ins}}$	0.031	-0.094	-0.259	0.077	0.262
	(0.047)	(0.087)	(0.055)	(0.035)	(0.049)
$\hat{\eta}_{kg(i)t}^{\text{Post}}$	-0.168	0.260	0.039	0.076	-0.012
	(0.057)	(0.124)	(0.016)	(0.034)	(0.021)
$\hat{\eta}$	0.018	0.145	-0.127	0.077	-0.113
	(0.039)	(0.039)	(0.013)	(0.018)	(0.023)
$\hat{\beta}$	0.535	0.131	0.255	0.218	0.515
	(0.011)	(0.039)	(0.025)	(0.027)	(0.028)
Observations	2924	404	2310	1484	4956
Δ Cost ($ per ref)	185	-259	235	-104	58
Referrals	656	127	576	304	1150

Table F.7: Outpatient Referral Steering - Commercial - No Random Effects

Notes: The table presents the model estimates for the change in the acquired PCP's cost sensitivity for inpatient referrals across outpatient specialist procedures for Commercial beneficiaries. An observation is a referral. We restrict our sample to our Acquired PCP sample, that is, we keep only patients of the acquired practice. The treated group are Commercial beneficiaries of the acquiring insurer at the acquired practice. The control group are Commercial beneficiaries of other insurers at the acquired practice. We do not include random effects in this specification.

	Evaluation and Management	
	Home Visit	Specialist Visit
$\hat{\eta}_{kg(i)t}^{\text{Acq Ins \& Post}}$	0.996	0.197
	(0.176)	(0.043)
$\hat{\eta}_{kg(i)t}^{\text{Acq Ins}}$	-0.971	0.156
	(0.136)	(0.056)
$\hat{\eta}_{kg(i)t}^{\text{Post}}$	-0.691	-0.187
	(0.349)	(0.032)
$\hat{\bar{\eta}}$	0.947	0.183
	(0.108)	(0.016)
$\hat{\beta}$	0.337	-0.424
	(0.066)	(0.044)
σ_δ	19.978	19.395
	(2.630)	(8.391)
Observations	3040	8182
Δ Cost ($ per ref)	784	227
Referrals	2552	5044

	Procedures				
	Major	Major Cardio	Major Ortho	Eye	Ambulatory
$\hat{\eta}_{kg(i)t}^{\text{Acq Ins \& Post}}$	0.548	-0.139	0.366	-0.014	0.036
	(1.580)	(0.436)	(0.123)	(0.030)	(0.138)
$\hat{\eta}_{kg(i)t}^{\text{Acq Ins}}$	-0.456	-0.282	-0.097	0.087	0.057
	(1.191)	(0.465)	(0.046)	(0.065)	(0.114)
$\hat{\eta}_{kg(i)t}^{\text{Post}}$	-0.430	0.569	-0.435	-0.066	0.041
	(1.284)	(0.166)	(0.069)	(0.062)	(0.118)
$\hat{\bar{\eta}}$	0.558	-0.129	0.254	0.109	0.069
	(0.999)	(0.152)	(0.068)	(0.020)	(0.158)
$\hat{\beta}$	0.559	0.464	0.385	0.116	0.435
	(0.041)	(0.288)	(0.171)	(0.014)	(0.068)
σ_δ	19.098	18.413	19.825	17.684	19.691
	(281.024)	(30.243)	(19.243)	(26.047)	(11.827)
Observations	955	428	1168	2312	2086
Δ Cost ($ per ref)	511	-138	369	-10	38
Referrals	639	287	761	1003	1371

Table F.8: Outpatient Referral Steering - MA - Random Effects

Notes: The table presents the model estimates for the change in the acquired PCP's cost sensitivity for inpatient referrals across outpatient specialist procedures for MA beneficiaries. An observation is a referral. We restrict our sample to our Acquired PCP sample, that is, we keep only patients of the acquired practice. The treated group are MA beneficiaries of the acquiring insurer at the acquired practice. The control group are MA beneficiaries of other insurers at the acquired practice. We include random effects in this specification.

	Evaluation and Management	
	Home Visit	Specialist Visit
$\hat{\eta}_{kg(i)t}^{\text{Acq Ins \& Post}}$	0.532	0.195
	(0.080)	(0.026)
$\hat{\eta}_{kg(i)t}^{\text{Acq Ins}}$	-0.168	-0.051
	(0.059)	(0.033)
$\hat{\eta}_{kg(i)t}^{\text{Post}}$	-0.442	-0.021
	(0.086)	(0.046)
$\hat{\bar{\eta}}$	0.213	0.186
	(0.052)	(0.018)
$\hat{\beta}$	-0.021	-0.364
	(0.016)	(0.042)
Observations	3040	8182
Δ Cost ($ per ref)	450	231
Referrals	2552	5044

	Procedures				
	Major	Major Cardio	Major Ortho	Eye	Ambulatory
$\hat{\eta}_{kg(i)t}^{\text{Acq Ins \& Post}}$	0.270	-0.151	-0.140	0.266	0.195
	(0.115)	(0.275)	(0.141)	(0.032)	(0.151)
$\hat{\eta}_{kg(i)t}^{\text{Acq Ins}}$	-0.266	-0.182	0.356	0.372	0.107
	(0.062)	(0.167)	(0.161)	(0.055)	(0.122)
$\hat{\eta}_{kg(i)t}^{\text{Post}}$	-0.296	0.437	-0.072	-0.043	-0.015
	(0.078)	(0.241)	(0.042)	(0.015)	(0.085)
$\hat{\bar{\eta}}$	0.260	0.042	-0.277	0.105	-0.096
	(0.045)	(0.075)	(0.056)	(0.020)	(0.069)
$\hat{\beta}$	0.501	0.073	0.040	0.044	0.472
	(0.042)	(0.028)	(0.062)	(0.022)	(0.036)
Observations	955	428	1168	2312	2086
Δ Cost ($ per ref)	243	-161	-138	153	184
Referrals	639	287	761	1003	1371

Table F.9: Outpatient Referral Steering - MA - No Random Effects

Notes: The table presents the model estimates for the change in the acquired PCP's cost sensitivity for inpatient referrals across outpatient specialist procedures for MA beneficiaries. An observation is a referral. We restrict our sample to our Acquired PCP sample, that is, we keep only patients of the acquired practice. The treated group are MA beneficiaries of the acquiring insurer at the acquired practice. The control group are MA beneficiaries of other insurers at the acquired practice. We do not include random effects in this specification.

	Appropriate			Not Appropriate		
	(1) MA	(2) Commercial	(3) All	(4) MA	(5) Commercial	(6) All
τ_{-2}	0.070	0.031	0.027	-0.040	-0.015	-0.004
	(0.074)	(0.015)	(0.013)	(0.093)	(0.018)	(0.017)
τ_{-1}	0	0	0	0	0	0
τ_0	0.007	0.029	0.027	-0.020	-0.029	-0.023
	(0.087)	(0.014)	(0.013)	(0.103)	(0.018)	(0.016)
τ_1	0.190	0.014	0.021	-0.137	0.011	-0.001
	(0.090)	(0.014)	(0.012)	(0.147)	(0.017)	(0.015)
τ_2	0.004	0.007	-0.000	-0.042	0.040	0.027
	(0.068)	(0.014)	(0.013)	(0.077)	(0.018)	(0.016)
Year FE	X	X	X	X	X	X
Baseline	28.90%	24.69%	25.71%	45.11%	49.35%	48.33%
Observations	5,967	22,769	28,736	5,967	22,769	28,736

Table F.10: Impact on Appropriateness of ER Visits

Notes: We construct a sample of ER visits and classify each visit using the Billings ED Algorithm. To construct the algorithm, a panel of ER and primary care physicians categorized each ER visit from a sample of approximately 6,000 ER records into one of four categories: (a) Not an emergency, e.g., could have waited 12 hours for care; (b) An emergency that a PCP instead of the ER could have treated ; (c) An emergency that could have been avoided, e.g., through preventative care; and (d) An emergency that could not have been avoided. Lastly, these classifications are mapped to the diagnoses provided in the visit, with each diagnosis having an assigned probability of each category. We utilize the diagnoses from the ER claims in our dataset and run them through the algorithm, resulting in a probabilistic assignment to each category for each visit. We further collapse these into Appropriate (d) and Not Appropriate (a, b, c), by summing over the probabilities in each category for easier interpretation. We conduct an event study design with individual and time fixed effects. In columns (1)-(3), we present the event study where we restrict to beneficiaries of the acquired practice and define treatment as beneficiaries of the acquiring insurer after the acquisition. In columns (4)-(6), we show the event study results for patients of all other PCPs, where we define treatment as adoption of a pay-per-patient contract with the acquiring insurer.

	Acquiring Insurer		Other Insurers	
	(1) MA	(2) Commercial	(3) MA	(4) Commercial
τ_{-2}	0.001	-0.002	-0.009	0.010
	(0.003)	(0.004)	(0.002)	(0.004)
τ_{-1}	0	0	0	0
τ_0	0.012	-0.019	-0.009	0.016
	(0.003)	(0.004)	(0.001)	(0.004)
τ_1	0.020	-0.033	-0.019	0.033
	(0.003)	(0.004)	(0.001)	(0.004)
τ_2	0.019	-0.029	-0.045	0.054
	(0.003)	(0.004)	(0.001)	(0.004)
Individual FE	X	X	X	X
Year FE	X	X	X	X
Baseline Share	11.42%	34.82%	3.36%	50.40%
Observations	3,686,171	3,686,171	3,686,171	3,686,171

Table F.11: Steering towards the Acquired Practice

Notes: This table presents the results of four event studies evaluating the change in likelihood of being a patient of the acquired PCP by group. Our four outcomes are indicators for beneficiaries' insurance plan: whether they are in the MA or Commercial segment, and whether they are beneficiaries of the acquiring insurer or another insurer.