Corporate Practice of Medicine: Vertical Alignment and Medicare Advantage Risk Coding

David J. Meyers, Jay Shroff, Yashaswini Singh, and Christopher Whaley*

November 2024

Abstract

A long-appreciated, but understudied, aspect of vertical consolidation and partnerships is the ability to align incentives to enable coordination between principals and agents. We examine this question in the context of vertical alignment between physicians and corporate primary care groups (CPCS), which are often integrated with Medicare Advantage (MA) insurers. While concerns of increased coding intensity in MA are well-documented, the mechanisms that enable aggressive coding are less clear. Using hand-collected data on physician employment by CPCPs and national MA encounter data, we find switching employment to a CPCP leads to an approximately 6% increase in MA risk score payments. We find no evidence of increased health care utilization and a *decrease* in an alternative risk score that uses pharmaceutical claims rather than recorded diagnoses. The increase in risk scores is driven by low-acuity and discretionary diagnoses and among patients in the lowest baseline risk score quartile. Higher risk-score payments lead to a $472 million increase in annual Medicare spending.

Keywords: Medicare Advantage, risk score payments, principle-agent relationships, health care provider consolidation

Preliminary and Incomplete

*All authors: Brown School of Public Health, Providence, RI.

An important economic impact of vertical integration between firms is the ability to coordinate actions and align incentives (Bresnahan and Levin 2012). Yet, examples of how coordination across originally separate firms occurs are limited. At the same time, vertical; integration, as with other forms of consolidation, can lead to increased consumer prices and market foreclosure (Ordover, Saloner, and Salop 1990). While most studies of vertical integration focus on product prices and firm output, many firms sell products to markets, particularly those with a large government presence, in which prices are administratively determined and less likely to be directly impacted by market forces.[1] In these cases, coordination across firms may be more important than in markets with direct prices.

A relevant and important example that we study comes from the Medicare program in the United States, which provides health insurance to nearly 70 million primarily elderly individuals using a hybrid coverage model. Approximately half of Medicare beneficiaries are enrolled in "Traditional" Medicare, in which health insurance benefits are directly administered by the federal government. The other half enroll in Medicare Advantage (MA), in which private insurers administer benefits, but are subsidized by the government in a risk-based payment arraignment. These risk payments are based on the medical diagnoses submitted by the patient's provider. However, a long-recognized challenge that the true illness severity is observed by neither the insurance plan nor the government (Cutler and Zeckhauser 2000).

Both the role of risk-based payments in insurance market stability and the potential for "gaming" of risk-based payments are the focus of numerous studies and policy proposals. However, the mechanisms behind risk score gaming have received less attention. Importantly, it

[1] For example, payment for government contractors is often determined by negotiations. Many nations use formulaically determine pharmaceutical prices.

is unclear why a physician who is responsible for medical coding would risk administrative time and hassle, reputational, and due to potential legal risks of billing fraud, financial costs of increased billing for the benefit of an insurer. This mechanism is critical, as the "agent" responsible for determining and treating a patient's diagnoses, physicians, are commonly separate from the "principal" that benefits from increased coding intensity, the insurer (McGuire 2000).

Recent years have seen fundamental changes to many aspects of the organization of physician markets that may enable coordination on risk-based coding. These organizational changes may align the incentives of physicians to more intensely code with insurers, who benefit from increased coding intensity. In particular, over the last decade, many physicians are now employed by corporate owners, including insurers and investment-backed organizations. These organizations are often closely linked with, and in many cases, directly owned by, Medicare Advantage insurance companies. Insurers and proponents commonly argue that this form of integration can increase care coordination and align incentives to promote "value-based" care. However, it can also align incentives to game risk-score payment measures. When independent, the physician operates as an agent primarily for the patient, but also for the Medicare program (i.e., the government) in determining and recording appropriate diagnoses. However, when employed by an insurer, there risks a potential that the physician's agency may shift to align with the incentives of the insurer when treating the patient and recording diagnoses. Importantly, while plans can provide some incentives to providers for additional risk score capture, integrating with the physician potentially allows for seamless coordination around risk scores to solve this coordination problem. Integrating with the physician potentially allows for coordination that solves this coordination problem.

In this paper, we examine these questions by measuring changes in risk coding intensity following a physician's change in employment to a corporate primary care practice (CPCP). To

do so, we leverage two unique sources of data. First, we collect unique data on physician employment by CPCP organizations. This process involves hand collect data on acquisitions of physician practices and using administrative data on provider employment. From 2016 to 2019, the number of physicians employed by a CPCP group in our sample increased from 87,000 to over 131,000. These integration events include both acquisitions by large publicly-traded insurers (e.g., UnitedHealth). We hand-collect data on primary care physician employment by CPCPs. We then link these data Medicare Advantage encounter data. These data include detailed information on health care utilization patterns. We use the data on beneficiary-level diagnoses to calculate the HCC risk score used by CMS for MA payments for each beneficiary-provider pairing.

With these data, we find that following acquisition by a CPCP group, the risk score for patients treated by a physician increases by approximately 6 percent. This result is robust to multiple specifications, including matching acquired physicians to observably similar non-acquired physicians and the inclusion of patient-level fixed effects. We also find no evidence of pre-acquisition trends in patient risk scores. In additional tests, we find several substantial heterogeneity changes in risk score coding. First, we find large differences across organizations. The largest increases are from Oakstreet Health, which has nearly a 40 percent increase in risk scores. We also find that risk scores are concentrated among beneficiaries with lower HCC scores at baseline.

Of course, one potential mechanism is that risk scores increase due to actual increased treatment, rather than simply increased diagnostic coding intensity. If increased risk scores are due to increased medical care, then patients may benefit through improved health. If, however, increased risk scores are solely due to increased coding intensity or documentation of ancillary codes, then patients won't benefit from increased care, and increased risk scores will represent a

transfer from the Medicare program to insurers. To test this mechanism, we conduct several tests of changes in medical care utilization. We find small changes, and slight increases, in either the number of physician visits, inpatient hospitalization days, or emergency department visits. Reductions in the latter two forms of care may reflect improved care coordination that is a supposed benefit of integration. We also find an approximately 15 percent increase in the number of HCC risk codes attributed a patient. This increase in diagnoses, which is not accompanied by changes in healthcare utilization indicative of higher health needs, suggest that the observed increase in risk scores is due to coding intensity. We further find that increased coding intensity is driven by lower-acuity conditions in which a physician has more discretion to diagnose, such as depression, rather than higher-acuity or more stark conditions such as advanced vascular disease.

Additionally, we test for changes in underlying health status using the Johns Hopkins ACG pharmaceutical risk score measure. Like other risk score measures, the ACG risk score is designed to predict future medical utilization and spending and has been used in other settings to measure patient health risks (Handel 2013) and to benchmark differences in MA risk inflation (Jacobs and Kronick 2021). Importantly, it is distinct from the HCC model used by CMS to determine payments to Medicare Advantage insurers. We use the version of the ACG based on just prescription drug claims, rather than medical diagnoses. If increased HCC scores are due to increases in patient illness burden, then it is likely that these same care needs would be reflected in increased pharmaceutical drug utilization. However, we find *decreases* in ACG risk scores of nearly 20 percent. This finding suggests that increases in HCC scores used to determine MA payment are driven by increased medical coding intensity, rather than increased patient illness.

Overall, our results point to a marked increase in risk scores following insurer acquisition of physician practices. Importantly, we do not find meaningful changes in several dimensions of

medical utilization or patient risk when using an alternative risk score measure. This increase in risk score leads to an approximately $720 per member per year increase in Medicare payments to insurers collectively leading to a $472 million increase in Medicare payments to insurers just among patients at our identified CPCPs. Our findings are consistent with vertical integration enabling coordination between insurers and physicians that would not have existed without this formal linkage.

Overall, our results contribute to several literatures. First, we contribute to a growing literature on the impacts of MA on health care spending, and whether MA plans are more efficient than Traditional Medicare. These studies find evidence of advantageous selection of healthier individuals, which accounts for modest increases in MA payments relative to comparable TM enrollees (Curto et al. 2019). More recent studies find evidence of increased payments to MA plans due to more aggressive coding (Meyers and Trivedi 2021; Murray et al. 2024). An important contributor to differences in payments between MA and FFS are potentially upcoded risk adjusted payments in Medicare Advantage, resulting in billions in higher public spending in the United States (Meyers and Trivedi 2021; Murrin 2021). However, these studies do not examine the underlying mechanisms that lead to increased coding intensity. Our study is most closely related to Geruso and Layton (2020), who exploit geographic shocks to MA plans and find increased coding intensity among MA plans, and even higher intensity among provider-owned MA. Our paper extends this research by providing a within-MA estimate of upcoding. We also bring improved data on provider affiliation from an emerging source of provider consolidation.

We also contribute to a more general literature on the role of risk scores in determining insurance payments. As summarized in Geruso and Layton (2017), much of the existing research on risk adjustment has focused on issues around the influence of risk scores on consumer or firm

choice coverage decisions and the role of risk adjustment in stabilizing insurance pools (Rothschild and Stiglitz 1978). However, few studies have examined the impacts of *ex-post* manipulation of risk scores by once plan selections are made. Finally, we contribute to the literature on physician agency, which has largely focused on treatment decisions made by providers (McGuire 2000; Gruber and Owings 1996). More recent studies examine the role of physician agency and financial incentives in determining choice of facility location (Munnich et al. 2020). We extend this literature by examining how insurer incentives can influence physician billing behavior and how consolidation enables alignment of physician and insurer incentives.

The rest of the paper proceeds as follows. We first provide an overview of Medicare Advantage and the recent emergence of corporate primary care practices. We then describe our data and empirical approach. Next, we present our main results, additional robustness tests and empirical tests that examine mechanisms. We then conclude.

Background on Medicare Advantage and Corporate Primary Care Practices

Growth in Medicare Advantage

Medicare Advantage (MA), an alternative to traditional Medicare, has experienced substantial growth over the last two decades. Enrollment in Medicare Advantage (MA)—private insurance plans for Medicare beneficiaries—has grown rapidly, from 25 percent of beneficiaries in 2010 to 54 percent in 2024.[2] Medicare Advantage (MA) plans are funded through a capitated payment system, meaning that the federal government pays private insurers a fixed, per-member, per-month

[2] Source: Kaiser Family Foundation https://www.kff.org/medicare/issue-brief/medicare-advantage-in-2024-enrollment-update-and-key-trends/

amount to cover each enrollee. This amount is intended to reflect the estimated costs of caring for individuals based on their health status and demographics, rather than on actual medical expenses. Payments to Medicare Advantage plans vary based on benchmarks set at the county level, adjusted annually, and modified by a risk adjustment model to more accurately align payments with enrollees' healthcare needs.

HCC model for risk score payments

The Hierarchical Condition Category (HCC) model is the primary tool Medicare uses to calculate risk scores and adjust payments for Medicare Advantage enrollees to ensure that plans with sicker populations receive higher payments. Developed by CMS, the HCC model assigns scores to each beneficiary based on their demographic information (age, gender, Medicaid eligibility) and documented medical conditions. Each condition correlates with an associated risk score or "weight" that reflects the expected cost of care. Higher risk scores indicate a greater expected cost, translating to higher payments to the MA plan.

A growing policy concern is that the HCC model creates financial incentives for MA insurers to aggressively code a beneficiaries' health status to make them appear to have additional illnesses and other conditions that would command higher payment. This occurrence, referred to as "upcoding," inflates risk scores and ultimately drive higher reimbursement to MA plans without beneficiaries receiving care for those diagnoses. The Office of Inspector General (OIG), CMS, the United States Department of Justice (DOJ), the Government Accountability Office (GAO), and the Medicare Payment Advisory Commission (MedPAC) have expressed concern related to MA companies' inflating their beneficiaries' risk scores. Research has also documented that MA companies differentially code more intensely than the Traditional Medicare program, resulting in

risk adjusted payments being a potential source of overpayments in the MA program (James et al. 2024; Jacobs and Kronick 2018).

Trends in PCP group integration

The growth in MA and related value-based programs has resulted in corporate interest in physician practices from private equity firms and health insurers aiming to provide investments to facilitate coordinated and comprehensive care. This trend is most notable within primary care, given primary care physicians hold information on patient medical history, informally coordinate care, and in some cases, formally manage access to specialists through referrals.

The growth of corporate primary care practices (CPCPs) has been supported predominantly by risk-based payment arrangements in MA. While data on new entrants in primary care is lacking, media reports suggest that CPCPs take the form of joint ventures between private equity firms and health insurers. For example, CPCPs include Oak Street Health, a network of primary care centers with investments from private equity firm General Atlantic and Humana that had an IPO in 2020 (eventually sold to CVS Health for $10.6 billion in 2023), Iora Health, a primary care platform with investments from PE investors and Humana (eventually sold to One Medical for $2.1 billion in 2021),[3] CenterWell Senior Primary Care (a joint venture of Humana with PE firm Welsh, Carson, Anderson & Stowe), OptumCare (an affiliate of UnitedHealth Group).

While a majority of CPCPs exclusively care for older adults in Medicare, particularly Medicare Advantage (and are thus marketed as "senior-focused primary care"), a smaller group of CPCPs target the commercial population. The commercial-focused CPCPs are comparable to concierge medicine, generally charging a flat per-employee-per-month fee to prioritize patient access to

[3] Additional information in Iora's expansion into MA populations is here: https://www.businessinsider.com/iora-health-raising-money-for-new-primary-care-model-2019-10

same-day appointments. An example of this trend is 1Life Healthcare (One Medical) that was acquired by Amazon for $3.9 billion in 2023).

Data

Corporate Primary Care Practice Employment Data

Despite the growth and interest in CPCPs, data on physicians who work for these organizations is not readily available. We compiled a hand-collected dataset on physician employment by CPCPs by linking novel affiliation data to the Medicare Data on Provider Practice and Specialty (MD-PPAS) from 2016-2020.

To do so, we first identified the target business name, along with variations of the legal business name (LBN), through internet searched, SEC filings, etc. We identified CPCPs that represent senior-focused primary care organizations. Next, we used partial or full string-matching techniques to link the TIN legal business name to the target business name. We limited the MD-PPAS data to only primary care physicians based on physician specialty codes (Family Practice, General Practice, Geriatric Medicine, Internal Medicine, Nurse Practitioner, OBGYN, Pediatric Medicine, Physician Assistant, Preventive Medicine).

Our treatment physician group is defined as physicians who switch affiliation from a non-CPCP TIN to a CPCP TIN from one year to the next, during 2017-2019 (Treatment-Switched). We also separately identify two control groups: first, we identify physicians who switch affiliation from a non-CPCP TIN to another non-CPCP TIN (Control-Switched), and second, we identify another group of physicians who do not switch TIN affiliation at any point during the study period (Control – Never Switched). Limit physicians to only PCPs (internal medicine), excluded nurse practitioners. The switches themselves likely take one of two forms- either an individual physician

who is choosing to change their enrollment from one practice or another, or cases when an entire practice is acquired by a CPCP entity and all or most affiliated physicians make a switch at once.

Medicare Advantage Encounter Data

We link the data on insurer-physician integration events to a 20% random national sample of Medicare Advantage encounter data over the 2016 to 2020 period. Similar to claims data, these data include detailed information on patient diagnoses, procedures performed, and provider information. However, unlike most claims data sources, the encounter data do not include information on prices negotiated between MA insurers and providers.

We use these data to construct our primary outcome measure of beneficiary-level HCC scores. The HCC scores account for all diagnoses received by a beneficiary. We construct several measures of beneficiary HCC scores. In our main analysis, we consider only the diagnoses coded by a given provider, identified using the provider's National Provider Identifier (NPI). To account for broader patterns of care and to examine differences based on baseline differences, we also measure a beneficiary's annual HCC score across all providers. The latter is used by CMS to determine payment to MA insurers. To calculate HCC scores, we use CMS Software that ingests diagnoses codes from the encounter data and calculates respective risk scores. The weights that each chronic condition attribute to a beneficiary's HCC score change over time. To avoid differences due to changes in the broader HCC model, we use the 2024 HCC weights across all years.[4] For additional outcomes, we also use the encounter data to calculate beneficiary-level measures of medical utilization, including hospitalizations.

[4] We use the 2024 HCC model: https://www.cms.gov/medicare/health-plans/medicareadvtgspecratestats/risk-adjustors/2024-model-software/icd-10-mappings

Empirical Approach

Our empirical approach is to estimate the effect of integration with a CPCP on physician risk score coding. To do so, we follow the approach used by several other studies that examine the impact of changes in physician markets and use a generalized difference-in-differences approach. In particular, we estimate event studies that estimate changes in HCC risk scores for patient i in year t matched to physician p:

$$HCC_{pt} = \alpha + \sum_{j} \beta_j \cdot D_{p,t-j} + \delta_t + \pi_p + X_{pt} + \varepsilon_{pt} \quad (1)$$

In this model, the set of β_j coefficients represent the differential change in HCC scores relative to the year of switch, D. We include fixed effects for year (δ_t), physician (π_p), as well as a set of physician-level time-varying controls (mean patient age, gender, share non-white). Due to concerns about staggered treatment timing and treatment effect heterogeneity, we estimate these event studies using a dynamic treatment effect estimator (Sun and Abraham 2020). For the "never switch" control group, we randomly assign a placebo switch date, similar to a "stacked" event study approach (Cengiz et al. 2019).

This approach relies on several assumptions. First, as with other difference-in-differences evaluations, we assume no contemporaneous shocks that occur at the time of provider switching. We test this assumption by examining pre-switch trends in integration.

A related concern is that provider switching is endogenous. In an additional test, we match switching providers to similar non-switching providers. Additionally, providers who switch may also change patient panels. We test for this effect in two ways. First, we estimate changes in

observable patient characteristics (e.g., patient age, gender, race) following switching. Second, we attribute beneficiaries to a PCP and estimate changes in *within-patient* HCC risk scores following provider switching.

Finally, our main outcome of interest is changes in HCC risk scores. However, these changes may be mechanically linked to increased health care utilization or diagnoses that follows provider switching. For example, new electronic health record platforms may not track a patient's previous diagnoses or treatments. We estimate changes in both overall health care utilization, as well as "high-value" care. To further test for changes in patient utilization, we use an alternative risk score, the ACG risk score model, which is based off of prescription drug utilization. If increased risk scores reflect higher health needs and utilization, then presumably these needs will also be reflected in medication use.

Results

Mean changes in HCC risk coding

Before describing our main empirical findings, Figure 2 presents unadjusted trends in the mean beneficiary HCC scores for our three physician cohorts–our treatment group of physicians who "switch" and are acquired by an insurer over this time period, our first control group of physicians who are "never treated" and are remain independent from insurers over our time period, and our secondary control group of physicians who switch affiliations during this period, but not to a CPCP. This figure shows no change in HCC scores among our physicians who remain with the same organization during the entire period, and a modest increase among those who change affiliation, but not to a CPCP. However, for the treatment cohort that switches to a CPCP, HCC scores increase following the switch. Without any regression adjustment, we find an approximately

6 percentage point increase in HCC scores. Because HCC scores are centered at 1.0, this unadjusted trend equates to an approximately 6 percent increase in HCC scores among physicians who vertically integrate with an insurer.

Figure 3 presents our main event study results that measure changes in HCC risk scores following a physician practice's vertical integration with an insurer. Prior to integration, we see small changes in an attributed patient's risk score. Following a physician switching, there is an immediate and steady increase in HCC risk scores. By the third-year post-switch, we see a 0.06-unit increase in HCC scores. Because HCC scores are normalized to mean 1.0, the observed increase translates to an approximately 6 percent increase in risk scores following a physician's switch to a CPCP.

To place these results in context, the "baseline" payment upon which Medicare Advantage payments are based is approximately $12,000. Under the MA payment model, a patient's HCC risk score is multiplied by the baseline payment rate to determine payment from CMS to the patient's MA plan. Thus, the 6 percent increase in HCC risk score translates to a $720 mean increase in CMS payments to MA insurers per beneficiary. In our analytic sample, a total of 131,000 MA beneficiaries are attributed to an integrated provider. Because our analytic sample is derived from a random 20% sample of encounter data, the national number of impacted beneficiaries is approximately 655,000. Thus, our results imply an annual $472 million increase in Medicare spending among patients receiving care from CPCPs alone. This increase in spending translates to about 0.1 percent of total annual Medicare Advantage spending.

A concern with our main approach is that patients may switch to a different physician once their PCP integrates with a CPCP practice. Alternatively, switching to a CPCP may expand the set of patients seen by a physician. To test the strength of our results to this concern, we estimate the

same regression model, but at the individual beneficiary level, rather than the physician level. Doing so allows us to include beneficiary fixed effects. Thus, these results estimate the *within-patient* change in risk score following the switch to a CPCP of a patient's primary care physician. As seen in Figure 4, the change in analysis to the beneficiary level and the inclusion of beneficiary fixed effects does not change our results. We still obtain a similar, 6 percentage point, increase in HCC scores by the third year of integration. The pre-switch trends in beneficiary HCC scores are flat.

Heterogeneous changes in HCC risk scores

Our main results focus on mean changes in risk scores following vertical integration. However, these mean results may mask important layers of heterogeneity. To examine heterogeneous effects, we implement two approaches. First, we apply the same empirical approach as Equation 1, but estimate separate regressions for each CPCP organization. This test allows us to test for organizational heterogeneity in risk coding intensity. The results in Figure 5 show a wide range of treatment effects across firms. On the low end, we observe a less than 0.5 percentage point increase HCCs for physicians that switch to an Optum-owned practice or Atrius.[5] At the other end, we see an approximately 35 percentage point increase in HCCs for Oak Street Health. Further highlighting the potential role of CPCPs enabling risk score coding coordination with insurers, following its 2020 IPO, Oak Street was sold to CVS/Aetna for nearly $11 billion in 2023.

Our second approach examines heterogeneity based on differences in patient illness burdens. While our main results show a mean HCC score increase of approximately 6 percent, it is not immediately clear how this mean increase is driven. From a policy perspective, finding an

[5] Atrius was acquired by Optum in May 2022, which is beyond the period of our study sample.

approximately across-the-board increase in risk scores (e.g., all beneficiaries increase by 10 percent) has substantially different implications and potential responses than finding a larger and more concentrated increase among a smaller population subset. Even in the concentrated score increase, it is immediately obvious where in the distribution of patient illness burden the risk score increases.

To examine patient-level heterogeneity, we examine differences in changes in beneficiary-level HCC scores following a beneficiary's primary care provider switching to a CPCP based on the pre-switch distribution of baseline HCC scores. We use the year prior to switching (and the placebo switch year for the main control group) to establish baseline HCC score and divide the distribution of beneficiary HCC scores into four quartiles. We estimate separate event studies (equation 1) for each quartile of baseline HCC scores.

As seen in Figure 6, we find no change in HCC scores for beneficiaries with the highest illness burdens, those in the top quartile of HCC scores. The entirety of the changes in HCC scores following a physician's switch to a CPCP is among beneficiaries in the bottom three quartiles, with the largest increase among the second and third quartiles. Thus, it does not appear that increases in HCC scores are among beneficiaries who already have high HCC scores. Instead, the change is at the lower-end of the HCC distribution and increasing coding intensity of those with lower baseline HCC scores.

Robustness Tests

Changes in patient panel characteristics

As discussed above, the validity of our core results has several potential threats, which we attempt to address in a series of sensitivity tests. A primary concern is changes to a physician's patient panel following integration. This concern is particularly relevant given the incentive for insurers to "cherry pick" and select healthier patient populations. Existing studies document how insurer networks contribute to advantageous risk selection (Shepard 2022). We estimate similar event studies, but replace our HCC outcome with measures of patient observable characteristics. As shown in Figure 7, we see no change in attributed patient age, race, or gender following integration. Thus, it does not appear that physician vertical integration with insurers changes the composition of treated patients, and thereby contributes to changes in patient risk scores.

Changes in healthcare utilization

Our primary results show increases in HCC risk scores of approximately 6 percent following physician-insurer vertical integration. However, HCC risk scores are determined by underlying patient billing records. If providers more aggressively document billing or potentially inappropriately assign diagnosis, then risk score increases may be a case of "upcoding" or fraudulent billing behavior. If, however, risk scores are reflective of increased patient care, then the welfare effects are less clear. Patients may benefit from improved medical care, even if Medicare spending and payments to MA insurers increases. At the same time, increases in "low-value" care could indicate that higher care management, a purported aim of CPCPs, is not being met.

We examine these mechanisms by estimating the effect of insurer vertical integration on several measures of medical utilization, including emergency room visits and inpatient

hospitalizations. Both measures are indicative of higher health needs and utilization. Figure 8 shows noisy, but slight increases in both forms of utilization.

Changes in alternative risk score measure

We also examine changes in beneficiary risk using an alternative risk score measure, the ACG pharmaceutical risk score algorithm developed by Johns Hopkins to measure patient-level health care utilization and spending. As discussed above, this risk score measure uses pharmaceutical claims data, rather than diagnosis codes, to estimate future patient health care spending and utilization. Conceptually, higher health needs reflected in increased HCC scores should also lead to increased prescription drug utilization, particularly for chronic conditions.

However, we find decreases in ACG risk scores following a switch to a CPCP (Figure 9). When estimated at the physician level, the mean ACG score for a physician's panel decreases substantially, by approximately 20 percentage points. We find a similar effect when estimating changes in ACG scores at the individual beneficiary level. Like the HCC scores, the ACG score is centered at mean 1.0. Thus, our results imply an approximately 20 percent relative reduction in ACG scores. This finding is in stark contrast to the observed increase in HCC scores in our main analysis.

Mechanism: Increases in Number of Diagnosis Codes and Low-Acuity Diagnoses

We next explore the underlying mechanisms behind increases in HCC scores. We first estimate the beneficiary-level changes in the number HCC codes, rather than HCC risk scores. As described in the Appendix, HCC risk scores are determined through summing a beneficiary's HCC code

weights for individual HCC codes. Thus, this test examines changes in the count of a beneficiary's HCC codes following their primary care provider joining a CPCP. As reported in Figure 10, we find an approximately 0.18-unit increase in the number of HCC codes assigned to a beneficiary. The mean number of HCC codes in our sample is 1.18. Thus, the number of HCC codes that contribute to a patient's HCC risk score expands by approximately 15 percent.

While this test shows an increase in diagnosis intensity, it does not show where increases are coming from and the codes that contribute to increased diagnoses. Understanding these differences is important for both understanding the mechanisms of increased coding intensity, as well as informing policy responses. To do so, we separate HCC codes into condition types that vary by clinical severity (e.g., depression vs. severe cardiovascular conditions). These conditions also vary by the ability of regulators to monitor a patient's true diagnosis vs. those in which a treating physician has more leeway to diagnose. For example, mis-diagnosing a patient with a heart attack or cancer will be easier for regulators to detect than a mental health or low-acuity diagnosis. Despite the lower HCC weight for low-acuity and more discretionary diagnosis codes, there may be more pressure to increase these types of codes vs. more complex or less discretionary codes, which, despite higher HCC payment weights, are easier to detect inappropriate coding.

Consistent with this hypothesis, the results in Figure 11 show the increase in HCC codes is driven by lower-acuity and more discretionary codes. Particularly, we find the largest increases for vascular disease without complications and depression. We find increases, albeit smaller, for circulatory and pulmonary conditions. We find no changes in less discretionary conditions. Notably, we find no change in vascular disease with other chronic conditions. These results suggest the notable increase in the number of HCC risk scores, which leads to increases in HCC risk scores,

are driven by lower-acuity diagnosis that are less likely to be open to observable and conflicting opinions by regulators and coding monitors.

Discussion

We find that when providers switch their affiliation to CPCPs, primary care companies with strong vertical ties with Medicare Advantage plans, that there is a 6% increase in risk scores, which may contribute to nearly $500 million in additional Medicare spending for these practices alone. Results were robust to alternative specifications concentrated among patients with lower acuity to begin with, and not associated with other significant changes in patient composition, health status, or utilization.

Importantly, increases in risk scores associated with CPCPs does not necessarily mean that increased coding is fraudulent. Many patients, including those in our study population likely have under diagnoses of certain conditions. However, in vertically aligned arrangements, the coordination problem of coding capture is largely solved, allowing for potentially large increases in payments to some insurers and providers, that other less integrated providers and insurers are not able to compete with. Linking physicians to CPCPs that are closely affiliated to an insurer, allows the physician to operate as a more-aligned agent for their new principal, the MA insurer.

Currently, CMS has very little ability to regulate the role that CPCPs are playing in the MA program. There are no national datasets of CPCPs and vertically integrated providers, thus why we needed to hand code the data to conduct this work. There are also potential concerns over CMS's ability to identify if any of this coding behavior is fraudulent. Currently, CMS uses the RADV audit system to audit plan HCC scores. RADV audits include comparing risk codes

reported by plans to those that are included in a patient's medical record. In CPCP and other vertically integrated arrangements, there is likely less distance between the codes captured in a health record and a plan, limiting the ability for the audits to detect meaningful differences. This again creates challenges for the overall MA market, as some insurers have the means to engage in these behaviors while other insurers may not have the means to do so and are thus unable to compete.

Historically, providers were independent actors that served as agents to patients in the healthcare system. With the rapid expansion of CPCPs and vertical integration, providers may no longer have the same level of independence, raising concerns for the current Medicare payment system.

References

Bresnahan, Timothy, and Jonathan Levin. 2012. "21. Vertical Integration and Market Structure." In *21. Vertical Integration and Market Structure*, 853–90. Princeton University Press. https://doi.org/10.1515/9781400845354-023.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. "The Effect of Minimum Wages on Low-Wage Jobs*." *The Quarterly Journal of Economics* 134 (3): 1405–54. https://doi.org/10.1093/qje/qjz014.

Curto, Vilsa, Liran Einav, Amy Finkelstein, Jonathan Levin, and Jay Bhattacharya. 2019. "Health Care Spending and Utilization in Public and Private Medicare." *American Economic Journal. Applied Economics* 11 (2): 302–32. https://doi.org/10.1257/app.20170295.

Cutler, David M., and Richard J. Zeckhauser. 2000. "Chapter 11 The Anatomy of Health Insurance." In *Handbook of Health Economics*, edited by Anthony J. Culyer and Joseph P. Newhouse, 1, Part A:563–643. Elsevier. http://www.sciencedirect.com/science/article/pii/S1574006400801705.

Geruso, Michael, and Timothy Layton. 2020. "Upcoding: Evidence from Medicare on Squishy Risk Adjustment." *Journal of Political Economy* 128 (3): 984–1026. https://doi.org/10.1086/704756.

Geruso, Michael, and Timothy J. Layton. 2017. "Selection in Health Insurance Markets and Its Policy Remedies." *Journal of Economic Perspectives* 31 (4): 23–50. https://doi.org/10.1257/jep.31.4.23.

Gruber, Jonathan, and Maria Owings. 1996. "Physician Financial Incentives and Cesarean Section Delivery." *The RAND Journal of Economics* 27 (1): 99–123. https://doi.org/10.2307/2555794.

Handel, Benjamin R. 2013. "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts." *American Economic Review* 103 (7): 2643–82. https://doi.org/10.1257/aer.103.7.2643.

Jacobs, Paul D., and Richard Kronick. 2018. "Getting What We Pay For: How Do Risk-Based Payments to Medicare Advantage Plans Compare with Alternative Measures of Beneficiary Health Risk?" *Health Services Research* 53 (6): 4997–5015. https://doi.org/10.1111/1475-6773.12977.

———. 2021. "The Effects of Coding Intensity in Medicare Advantage on Plan Benefits and Finances." *Health Services Research* 56 (2): 178–87. https://doi.org/10.1111/1475-6773.13591.

James, Hannah O., Beth A. Dana, Momotazur Rahman, Daeho Kim, Amal N. Trivedi, Cyrus M. Kosar, and David J. Meyers. 2024. "Medicare Advantage Health Risk Assessments Contribute Up To $12 Billion Per Year To Risk-Adjusted Payments." *Health Affairs* 43 (5): 614–22. https://doi.org/10.1377/hlthaff.2023.00787.

McGuire, Thomas G. 2000. "Chapter 9 Physician Agency." In , edited by BT - Handbook of Health Economics, 1, Part A:461–536. Elsevier. http://www.sciencedirect.com/science/article/pii/S1574006400801687.

Meyers, David J., and Amal N. Trivedi. 2021. "Medicare Advantage Chart Reviews Are Associated with Billions in Additional Payments for Some Plans." *Medical Care* 59 (2): 96–100.

Munnich, Elizabeth L., Michael R. Richards, Christopher M. Whaley, and Xiaoxi Zhao. 2020. *Raising the Stakes: Physician Facility Investments and Provider Agency*. RAND Corporation. https://doi.org/10.7249/WRA621-4.

Murray, Roslyn C., David J. Meyers, Erin C Fuse Brown, Travis C. Williams, and Andrew M. Ryan. 2024. "Association Between the Medicare Advantage Quartile Adjustment System and Plan Behavior and Enrollment." *JAMA Health Forum* 5 (1): e234822. https://doi.org/10.1001/jamahealthforum.2023.4822.

Murrin, Suzanne. 2021. "Some Medicare Advantage Companies Leveraged Chart Reviews and Health Risk Assessments To Disproportionately Drive Payments." U.S. Department of Health and Human Services Office of Inspector General.

Ordover, Janusz A., Garth Saloner, and Steven C. Salop. 1990. "Equilibrium Vertical Foreclosure." *The American Economic Review* 80 (1): 127–42.

Rothschild, MICHAEL, and JOSEPH Stiglitz. 1978. "EQUILIBRIUM IN COMPETITIVE INSURANCE MARKETS: AN ESSAY ON THE ECONOMICS OF IMPERFECT INFORMATION*." In *Uncertainty in Economics*, edited by PETER Diamond and MICHAEL Rothschild, 257–80. Academic Press. https://doi.org/10.1016/B978-0-12-214850-7.50024-3.

Shepard, Mark. 2022. "Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange." *American Economic Review* 112 (2): 578–615. https://doi.org/10.1257/aer.20201453.

Sun, Liyang, and Sarah Abraham. 2020. "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects." *Journal of Econometrics*.

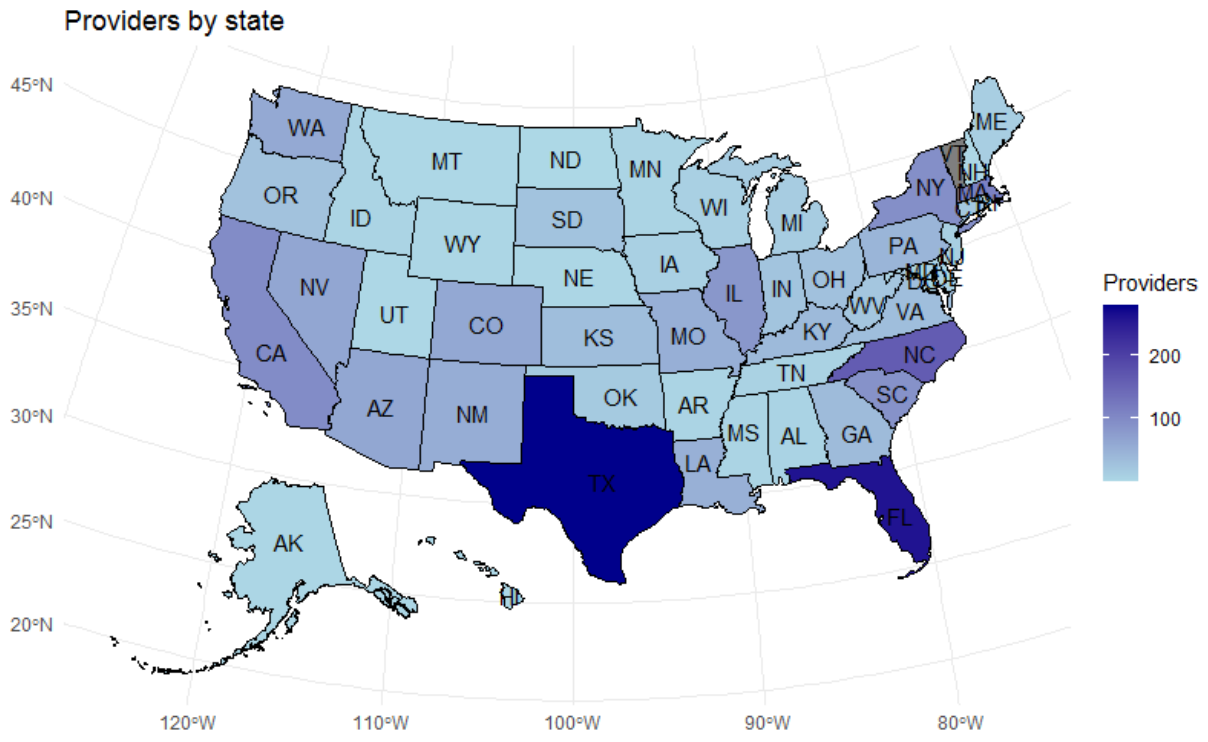
Tables and Figures

Table 1: Characteristics of Study Sample

	Control No Switch	Control Switched	Treatment
Degree			
MD	44.11	34.49	57.75
Missing	18.66	22.23	14.18
PA	12.92	15.86	10.58
DO	6.62	6.29	7.41
NP	6.60	7.73	3.60
Specialty			
Nurse Practitioner	27.28	33.60	14.51
Family Practice	21.74	18.18	33.43
Internal Medicine	19.85	14.63	24.54
Physician Assistant	17.01	21.56	13.74
Obstetrics & Gynecology	6.99	4.92	2.99
Missing	3.20	2.06	4.35
Hospitalist	1.66	3.50	4.94
General Practice	1.16	0.88	1.12

Figure 1: Geographic Distribution of CPCP Providers

(a) Number of CPCPs per State



(b) Trends in CPCP Growth per State

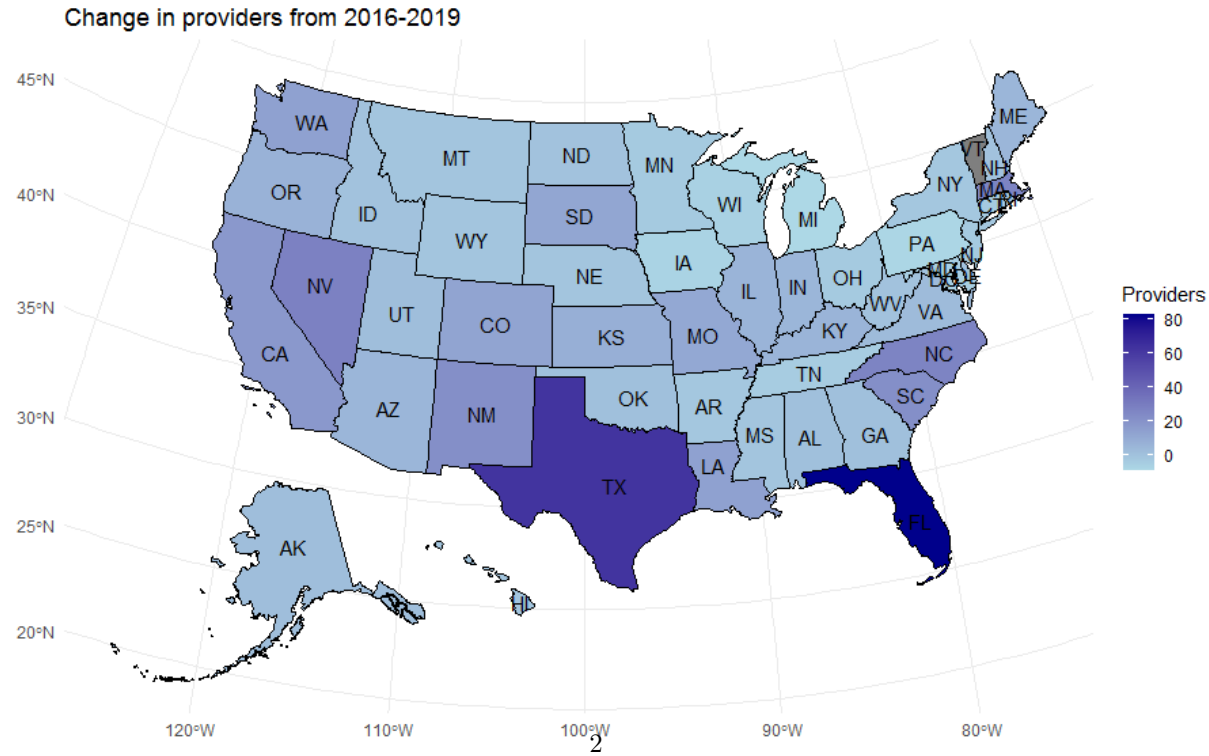


Figure 2: Unadjusted Trends in Provider Mean HCC Scores Relative to Switch to CPCP

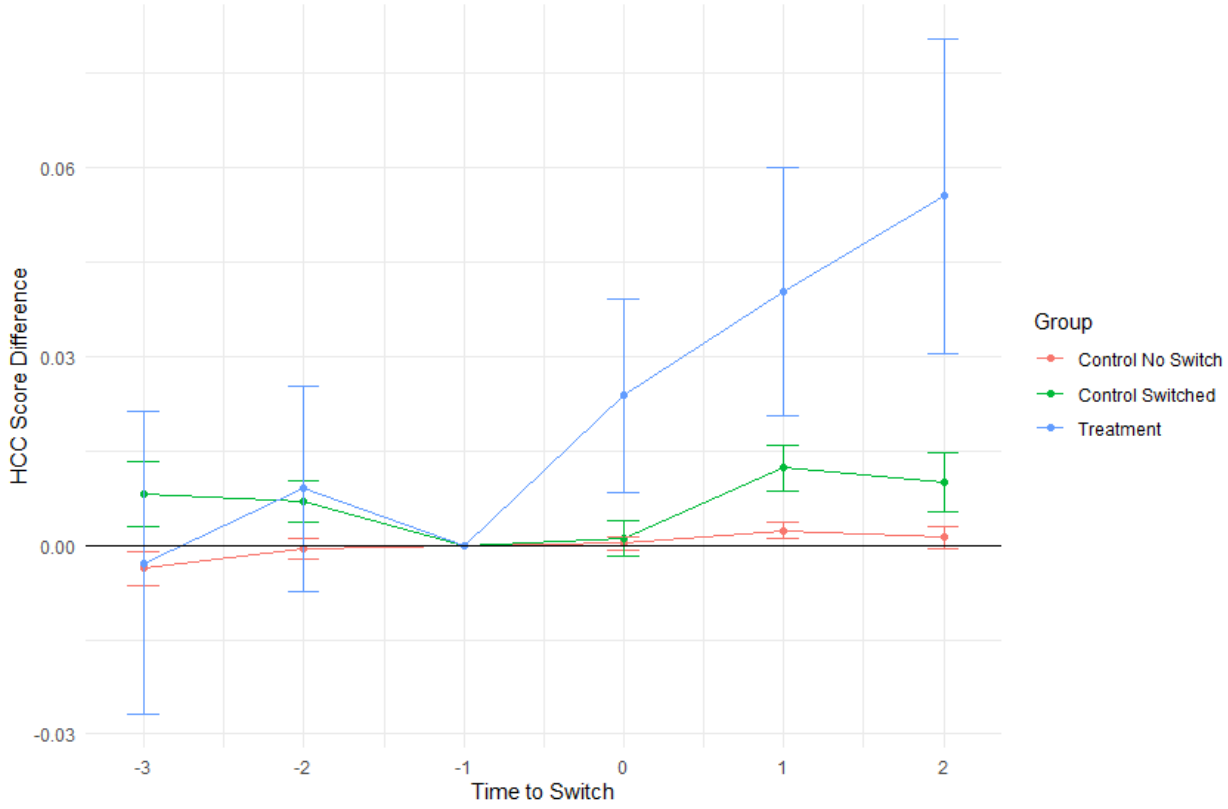


Figure 3: Change in Provider Mean HCC Scores Following Switch to CPCP

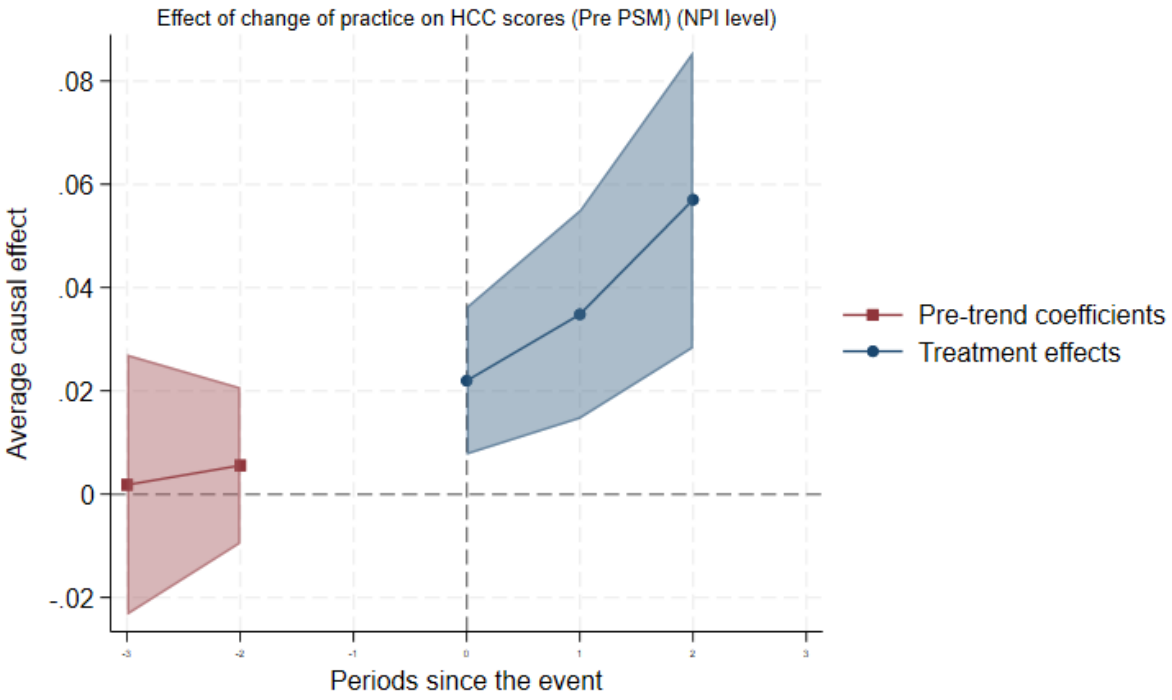


Figure 4: Change in Beneficiary HCC Scores Following PCP Switch

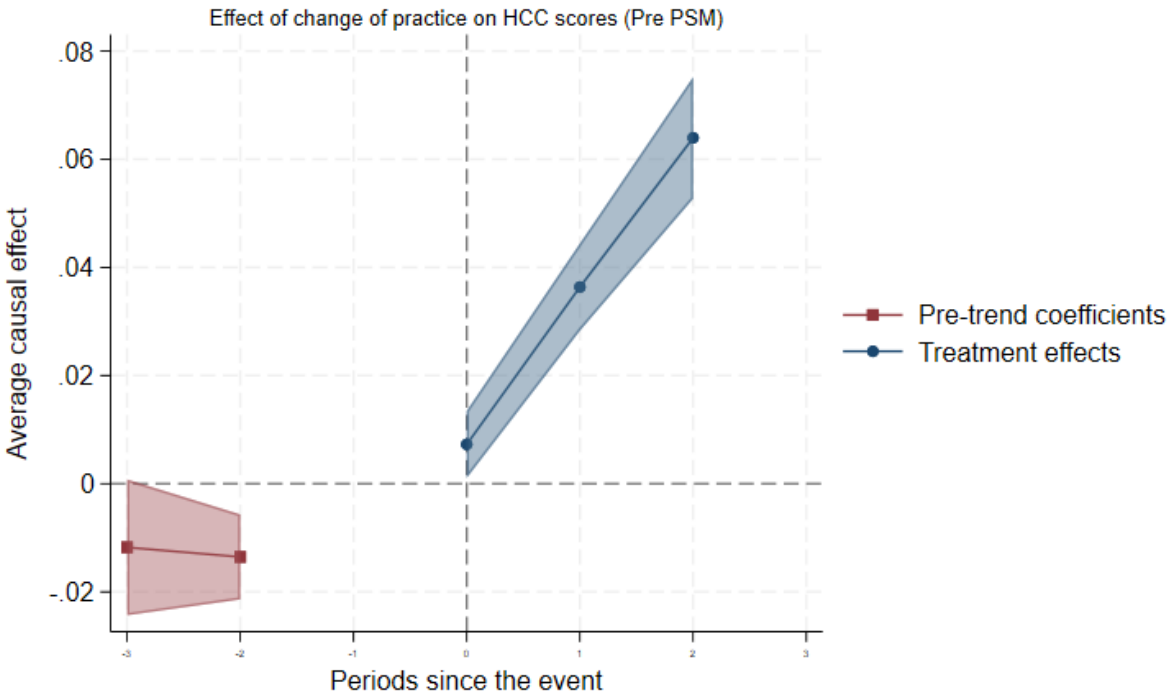


Figure 5: Firm-Level Change in Provider HCC Scores Following Switch to CPCP

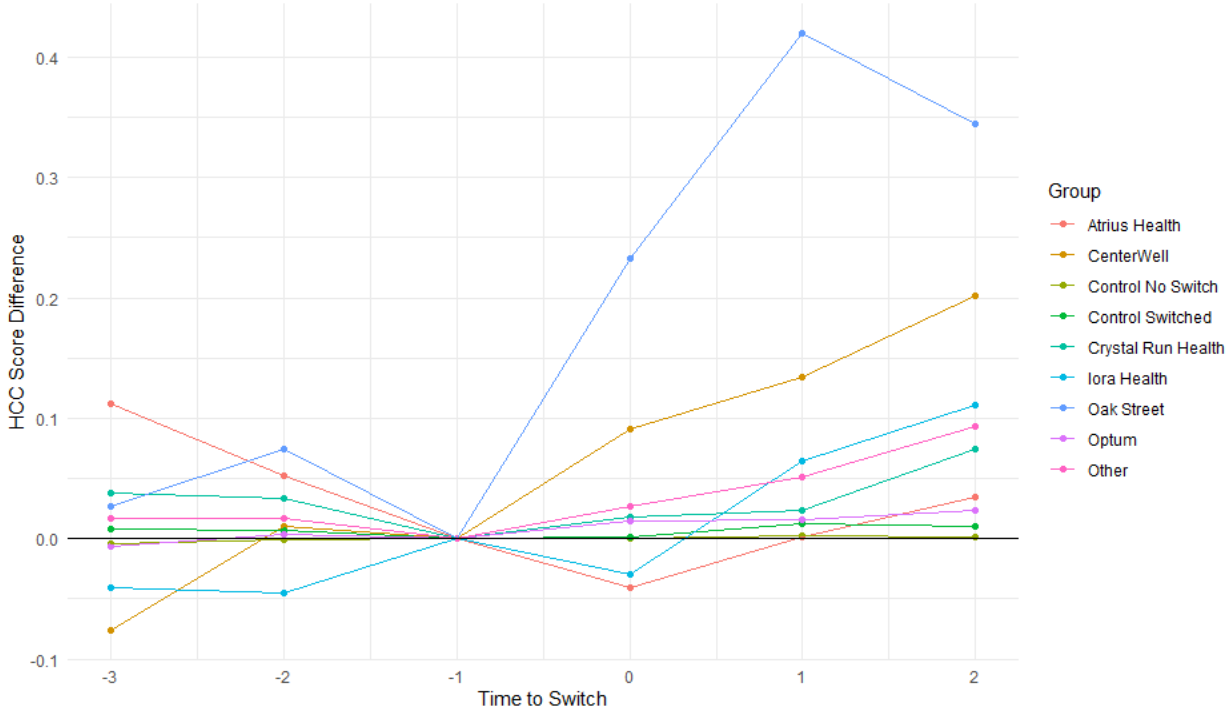


Figure 6: Beneficiary Heterogeneous Changes in HCC Scores Following Switch to CPCP

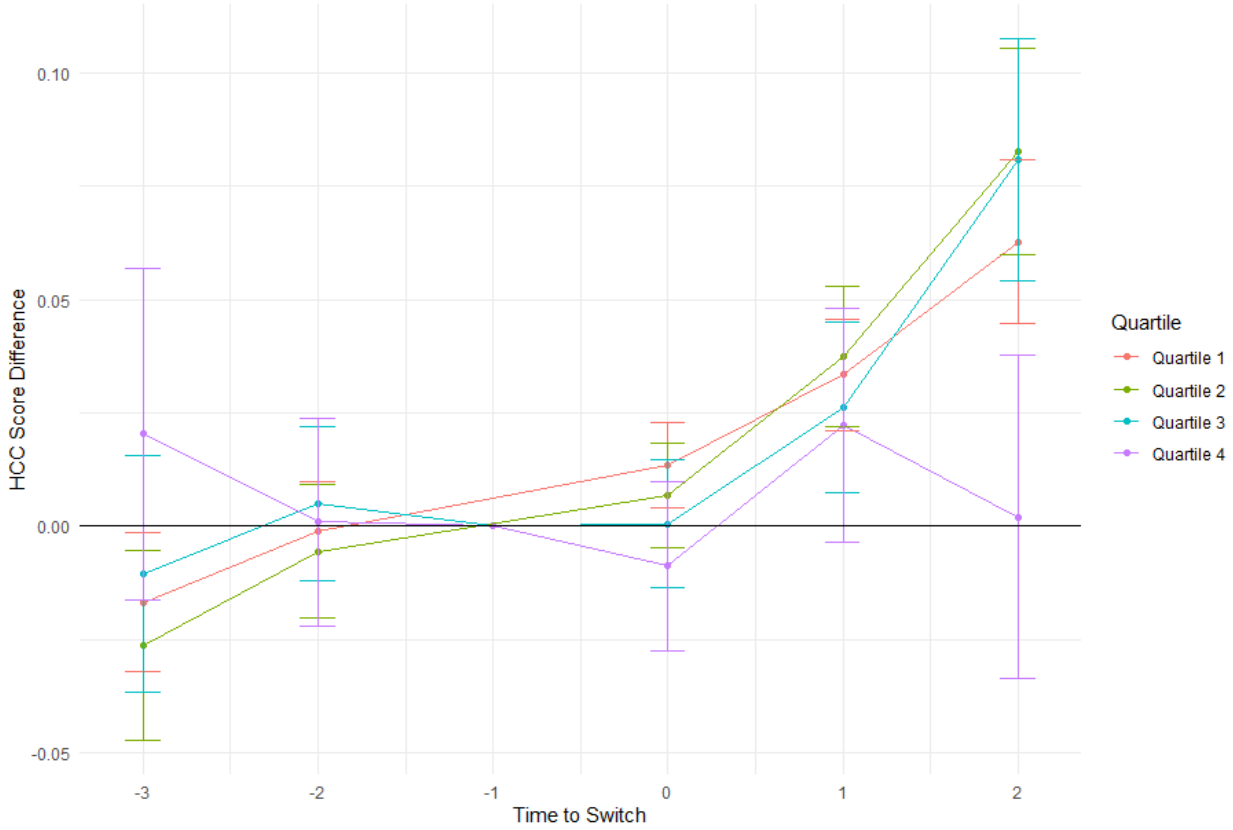
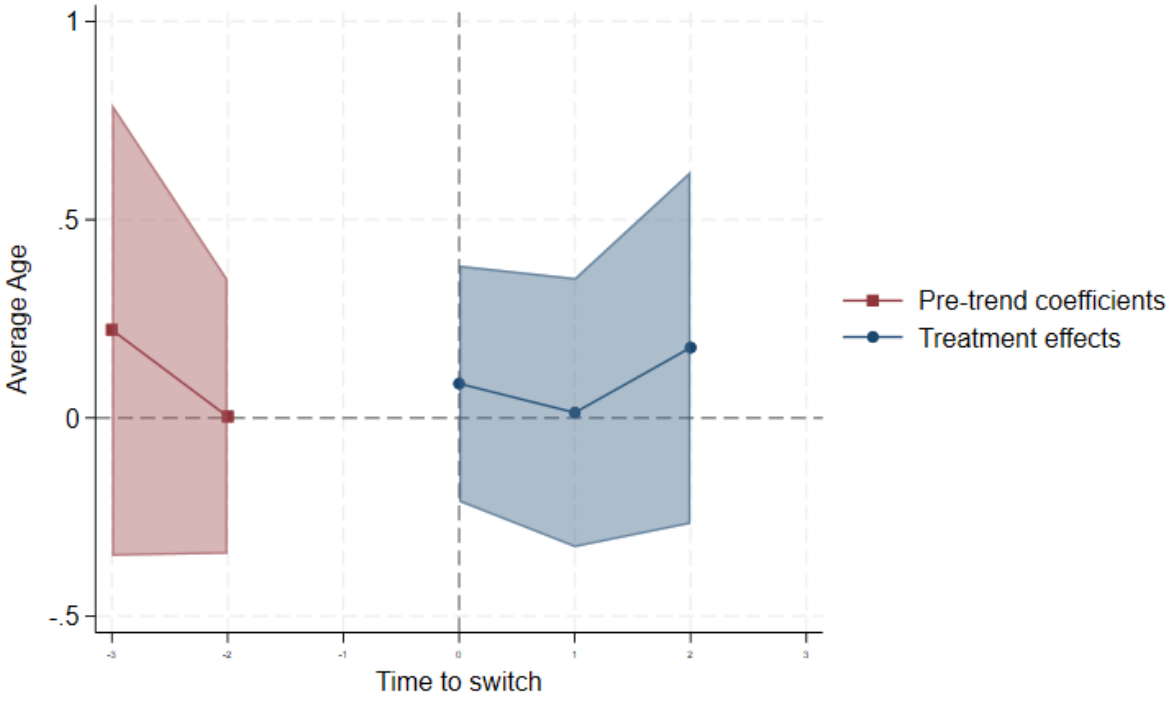


Figure 7: Change in Patient Characteristics Following Switch to CPCP

(a) Mean Patient Age



(b) Mean Patient Percent Male

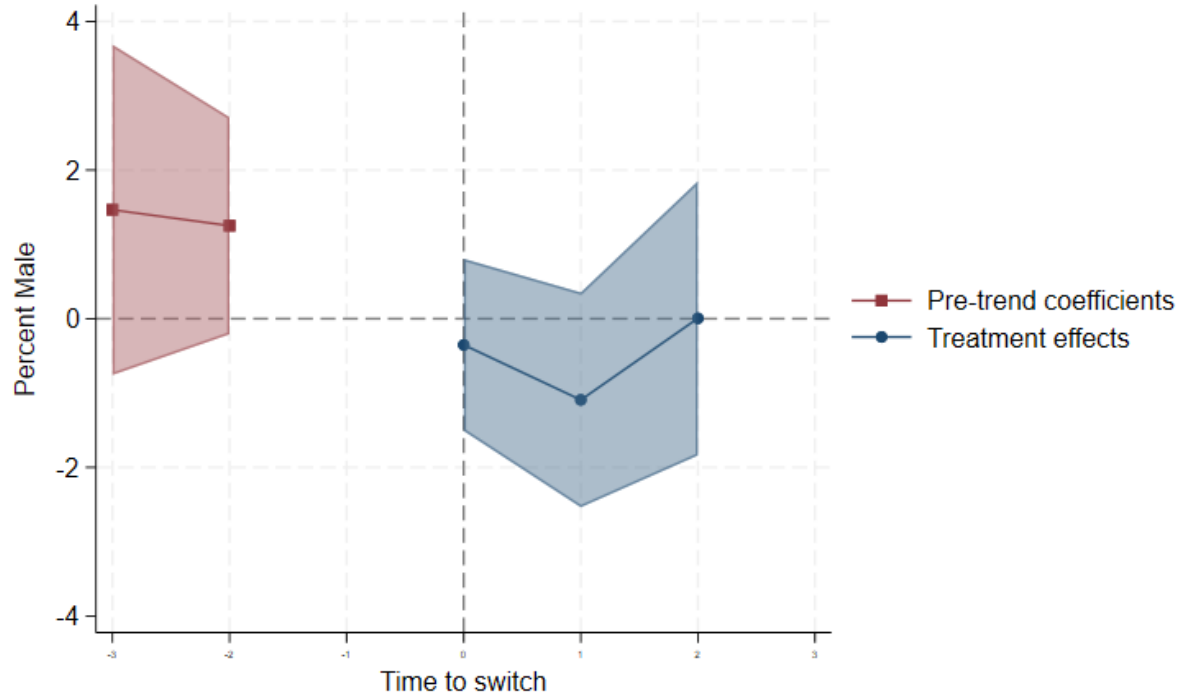
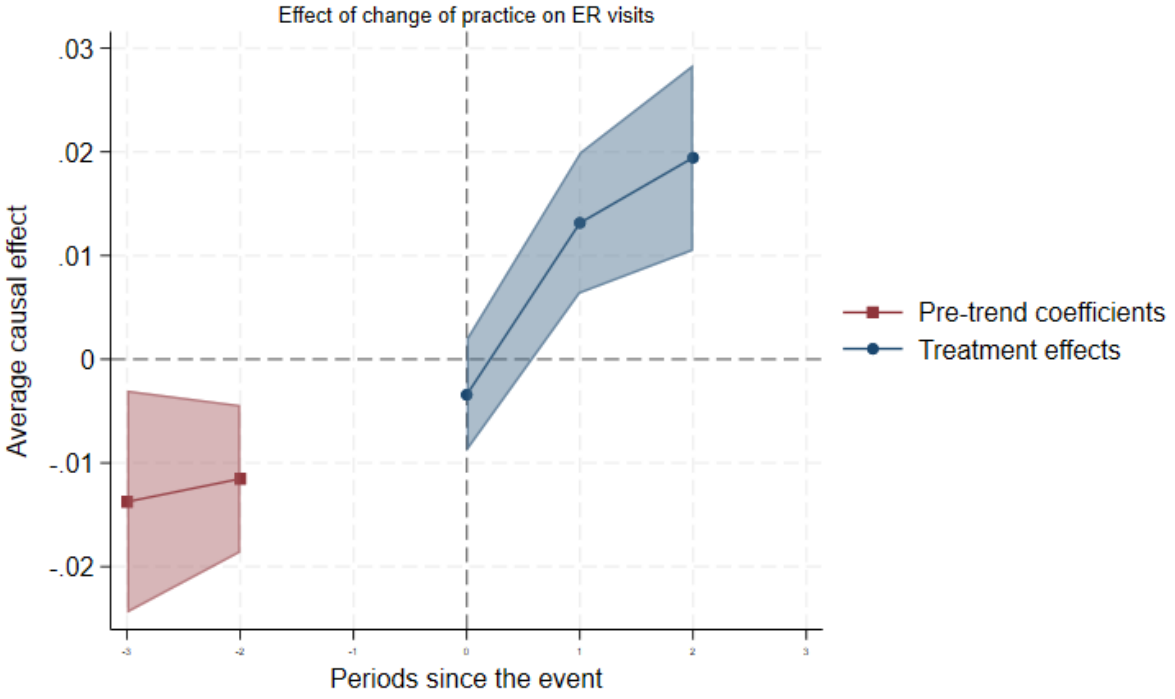


Figure 8: Change in Healthcare Utilization Following Switch to CPCP

(a) Emergency Room Visits



(b) Inpatient Visits

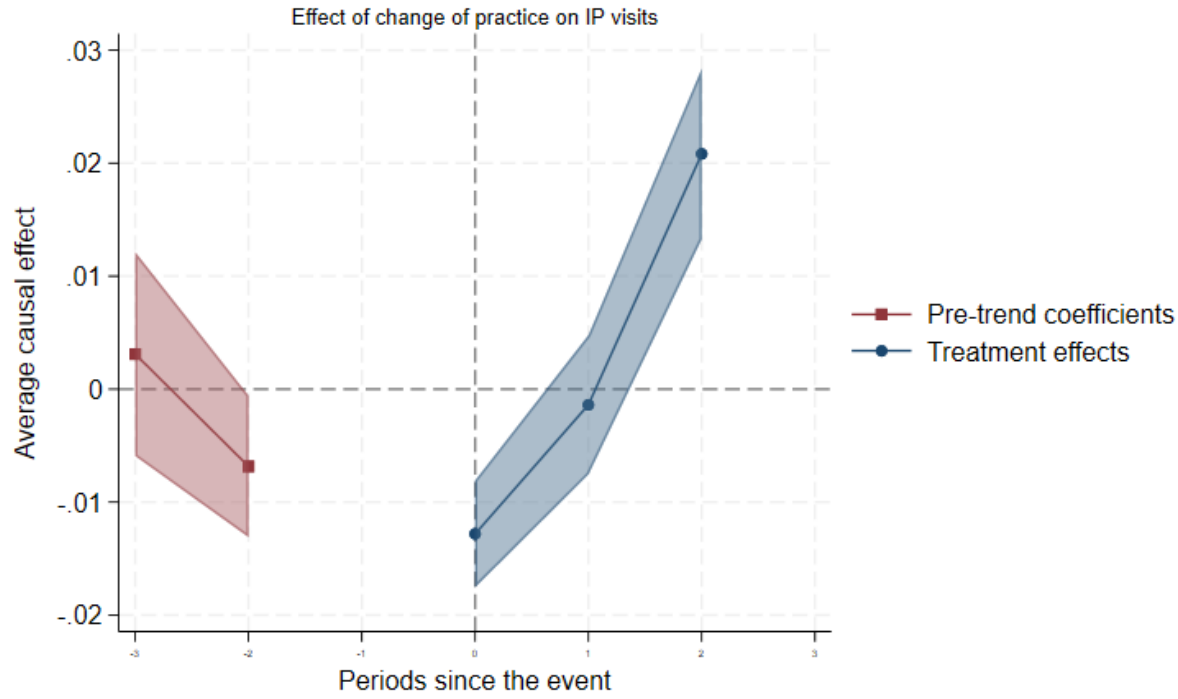
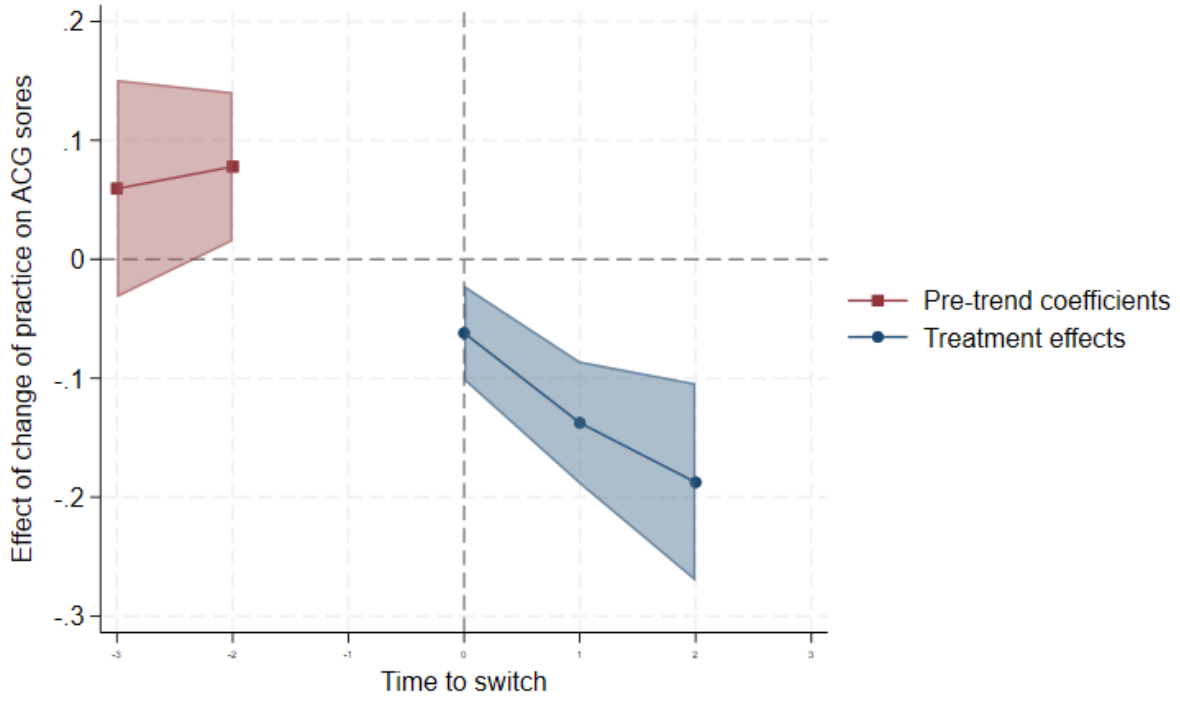


Figure 9: Change in ACG Risk Score Following Switch to CPCP

(a) Physician Level



(b) Beneficiary Level

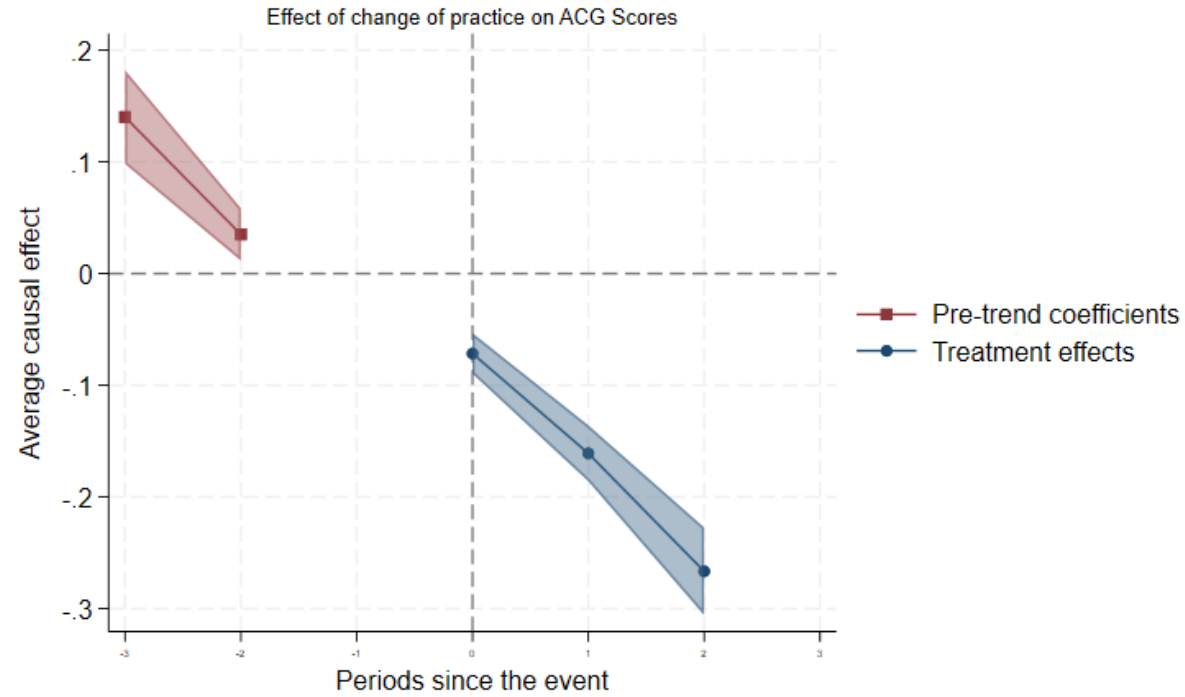


Figure 10: Changes in Number of HCC Codes Following Switch to CPCP

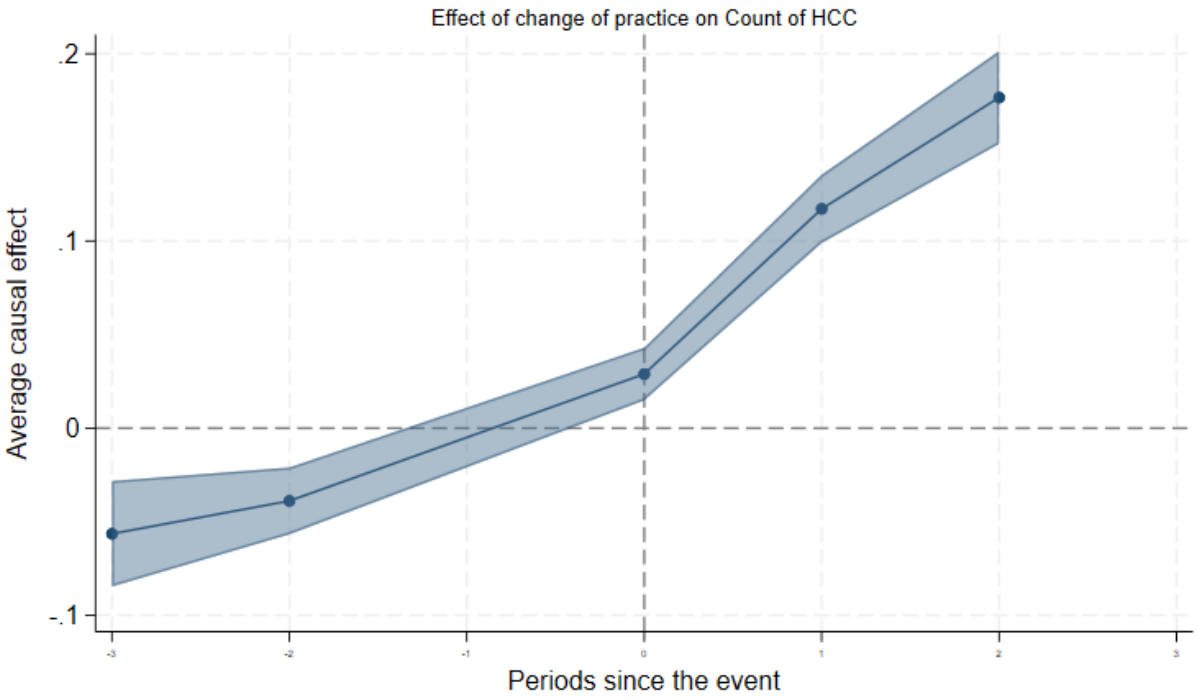


Figure 11: Changes in Specific HCC Categories Following Switch to CPCP

