Provider payment incentives: evidence from the U.S. hospice industry

Norma B. Coe[a] and David A. Rosenkranz[b]

December 10, 2024

(an updated version may be available here)

Abstract. Capping health care providers' average revenue can reduce allocative inefficiency under capitation. But its potential cost-savings may be undercut by health care providers who churn their patient censuses. We investigate this possibility in the U.S. hospice industry, where Medicare pays hospice programs fixed daily rates but caps their average annual revenue. By leveraging variation generated by the cap's nonlinear design and the transition between fiscal years, we find that programs on track to exceed the cap raise enrollment rates by 5.8% and live discharge rates by 4.3% in the fourth quarter, reducing financial penalties—but far from eliminating them. Marginal enrollees have longer average remaining lifetimes and more fragmented hospice spells, suggesting weaker intrinsic demand for hospice care. We discuss the cap's implications for market structure.

Keywords: Capitation; provider-induced demand; gaming; non-linear program design; Medicare; hospice

JEL No.: I0, I1, I11, I18

[a] Perelman School of Medicine at the University of Pennsylvania, 423 Guardian Drive, Blockley Hall, Philadelphia, PA 19104, USA. Email: nbcoe@pennmedicine.upenn.edu. Declarations of interest: none.

[b] Corresponding author. Department of Economics at Fordham University, 441 E Fordham Road, Dealy Hall, Bronx, NY 10458, USA. Email: drosenkranz@fordham.edu. Declarations of interest: none.

We thank Douglas Almond, Rebecca Anhang Price, Tyler Braun, Jeffrey Clemens, Natalie Ernecoff, Len Goff, Steven Joffe, Ryan McDevitt, Katherine Miller, Sungchul Park, Marc Remer, Heather Royer, Harald Schmidt, Sally Stearns, Chuxuan Sun, Ashley Swanson, Joan Teno, Lindsay White, three anonymous referees, and seminar participants at ASHEcon, AcademyHealth ARM, Bates College, Bowdoin College, Fordham University, the Maine Economics Conference, Rutgers University, Swarthmore College, the University of Pennsylvania, and WEAI. We also thank Kenneth Albert, Ted Cuppett, Theresa Forster, and others at the National Association for Home Care & Hospice for sharing their perspectives. This research is supported by the National Institutes of Health's National Institute on Aging (grant no. R01AG062595). Rosenkranz was supported by the Agency for Healthcare Research Quality National Research Service Award (T32), through grant number 5T32HS026116. We obtained IRB approval for this study through the University of Pennsylvania, protocol #831952. The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

1 Introduction

Concerns about high and rising health care costs have driven policymaking for decades. Inefficient health care spending can arise from asymmetric information and moral hazard, but well designed payment models can correct these market failures by shifting unnecessary costs onto providers (Ellis and McGuire 1986; Arrow 1963). For instance, under capitation, providers are paid fixed fees per unit of time (Friedberg et al. 2015). They therefore bear the marginal cost of care intensity and may lower costs by eliminating productive inefficiencies. But if they also treat patients longer than necessary, then capitation may nevertheless raise costs by creating new allocative inefficiencies. Can payers cut costs by controlling both productive and allocative inefficiencies simultaneously?

We investigate this question in the U.S. hospice industry, where hospice programs provide palliative care for terminally ill individuals and Medicare aims to reduce spending by combining capitation with a cap on their average annual revenue. The cap creates a maximum regulated revenue (MRR) for hospice programs given their patient volume. In theory, it can cut costs arising from unnecessarily long hospice stays because excess revenue is a cap liability that must be repaid. We show that hospice programs can undercut the cap by churning patients: a discharge can reduce revenue and an enrollment can increase the MRR by tens of thousands of dollars. Since enrollments can reduce cap liabilities, the cap may increase—not decrease—Medicare hospice spending. It also creates a unique opportunity to study provider-induced demand for end-of-life (EOL) care.

We identify churning in 20 years of hospice claims data using variation generated by the cap's nonlinear design and the transition between fiscal years. First, we observe that hospice programs on track to exceed the cap have a financial incentive to churn patients that other programs do not. Second, we observe that this incentive decreases sharply during the transition between fiscal years because cap liabilities are determined annually. We combine these observations in a difference-in-differences (DID) design that compares hospice programs on track to exceed the cap to otherwise similar programs not on track to exceed the cap during the transition between fiscal years.

We find that hospice programs on track to exceed the cap in a given fiscal year raised enrollment rates by 5.8% and live discharge rates by 4.3% on average in the fiscal year's fourth quarter. The temporal findings are telling: we find that their weekly enrollment rates differentially increase throughout the last quarter of the outgoing fiscal year before suddenly differentially decreasing by

13.9% on average in the first week of the next fiscal year. But the magnitude of this churning is small. It amounts to an additional 1.54 enrollments and 0.37 live discharges per program-year, or $30,000-$50,00 lower cap liabilities at most for an average hospice program. By contrast, we find that the 11% of program-years that exceeded the cap repaid $500,000 on average (or 21% of their gross Medicare revenue). We also find that the marginal enrollees were less likely to have been recently hospitalized or ultimately die in hospice care, suggesting that the cap caused programs to move down the demand curve and enroll patients with a lower intrinsic demand for hospice care.

Our design overcomes two identification problems that arise in within-program or between-program comparisons alone. First, programs that exceed the cap differ from others in several ways. For instance, consistent with existing evidence, we find that they are more likely to be newer, smaller, and for-profit. They also treat patients with higher lifetime lengths-of-stay in hospice care (LLOS), higher rates of neurological and cardiovascular conditions, and lower rates of cancer (MedPAC 2022). Second, there is seasonality in hospice utilization. For instance, we find that enrollment rates are higher on average during the winter and lower on average during holidays. Our DID analysis overcomes these identification problems by comparing outcome trends during the transition between fiscal years among programs on track to be above versus below the cap.

Our findings contribute to three strands of literature in economics and health policy. First, there is a large literature studying cost containment policies, such as deductibles (e.g., Brot-Goldberg et al. 2017), prospective payments (e.g., Grabowski et al. 2011; Meltzer et al. 2002; Pauly 2000), and certificate-of-need programs (e.g., Rosenkranz 2024; Polsky et al. 2014), among others (e.g., Alexander 2020; Park et al. 2017; Ho and Pakes 2014). Relative to a patient-level revenue cap, an average revenue cap does not penalize providers for occasionally treating patients with unexpectedly long stays. However, it creates a purely financial mechanism for one patient's care to affect another's. We measure the extent of this externality in hospice, where studies have shown that live discharge can be a costly disruption to care continuity (e.g., Dolin et al. 2017). Our findings suggest that despite creating an opportunity for churning—which providers exercised to some extent—an average revenue cap can reduce Medicare spending under capitation.[1]

Second, several studies examine how providers respond to financial incentives, such as those

[1]We focus on spending in the market where the cap exists—hospice care. Other studies have shown that hospice care can lower health care spending in other markets, in part because it is a substitute for curative care (e.g., Gruber et al. 2023; Zuckerman et al. 2015; Kelley, et al. 2013).

generated by reimbursement rates (e.g., Alexander and Schnell 2024; Clemens and Gottlieb 2014), bundling (e.g., Eliason et al. 2022; Einav et al. 2022), pay-for-performance programs (e.g., Gupta 2021; Eliason et al. 2018; Einav et al. 2018; Norton et al. 2018), and other factors (Chandra et al. 2012). Several studies have documented how insurers and providers may sometimes "game" payment systems to maximize profit (e.g., Gupta et al. 2024; Decarolis 2015; Dafny 2005) Outside health economics, Liebman and Mahoney (2017) found that expiring federal budgets are associated with higher year-end government spending. Similarly, we find that an annual revenue cap causes hospice programs to churn patients near the end of the fiscal year. But our findings suggest that their scope for inducing demand may be smaller than policymakers had feared. We hypothesize that programs are limited by non-pecuniary features of the Medicare hospice benefit, including that a non-hospice provider must certify that a prospective enrollee's life expectancy is 6 months or less and that hospice enrollees forgo Medicare coverage for their terminal illnesses.

Third, several studies describe the economics of the U.S. hospice industry and the cap. There is a positive correlation between cap liabilities and live discharge rates (e.g., Dolin et al. 2018; Plotzke et al. 2015; Teno et al. 2014), and surveys and press reports have identified instances where staff felt pressured to "pad the roster with new patients" or discharge patients with "overly long stays" (Kofman 2022; Jenkins et al. 2011; Sack 2007). But to date, it is unclear how widespread these practices are or whether the correlations are attributable to the cap versus other differences between hospice programs. We use the universe of Medicare hospice claims and a novel identification strategy to show that the cap causes hospice programs to raise both enrollment and live discharge rates, and to characterize the marginally enrolled and live discharged populations.

Finally, the Medicare Payment Advisory Commission (MedPAC) recently proposed reducing the cap and pegging it to geographic variation in hospice payment rates (MedPAC 2020). We find no evidence of bunching in the distribution of average annual revenue and a positive association between cap liabilities and reimburesement rates, suggesting that MedPAC's proposal could further reduce Medicare hospice spending and "mak[e] the cap more equitable across providers" (MedPAC 2020). However, we also find a positive association between cap liabilities and exit, which raises questions about the cap's potential effect on market structure and quality (Ata et al. 2012).

Our work is related to Gruber et al. (2023), who measure the association between a hospice program's probability of exceeding the cap in the fiscal year of a given month m and the probability

that a patient enrolled there in month m is live discharged within twelve months. Their regression analysis resembles a two-way fixed effects (TWFE) DID regression comparing live discharge trends between programs that experience larger versus smaller increases in cap risk within a fiscal year. However, causal inference is complicated by the possibility that live discharge shocks may affect both cap risk and live discharge trends. By contrast, we identify cap-induced live discharges using the transition between fiscal years, when programs on track to exceed the cap in an outgoing fiscal year experience a sharp change in their cap-induced financial incentives. We also study the cap's effect on other outcomes (e.g., enrollment) and its relationship to reimbursement rates and exit, whereas they separately study the effect of hospice care on total health care spending.

The remainder of this paper proceeds as follows. In section 2, we describe this study's economic and policy context. In section 3, we describe the cap-induced financial incentives. In section 4, we describe our data. In section 5, we report descriptive statistics. In section 6, we present our DID analysis. In section 7, we conclude.

2 Context

2.1 Hospice care and the U.S. hospice industry

Hospice care is a bundle of EOL health care services available to all terminally ill Medicare beneficiaries with a predicted life expectancy of six months or less. Its purpose is to improve a person's quality-of-life with pain and symptom relief, and keep them home with family and friends rather than an inpatient care setting. It often includes select medical services, homemaker services, and grief counseling. Medicare beneficiaries who enroll in hospice care forgo Medicare coverage for all other services related to their terminal illness for the duration of their hospice enrollment, including curative care.[2] Live discharge from hospice care may occur if the enrollee's condition improves or they choose to resume curative care. When available, a non-hospice physician involved in the beneficiary's care must be consulted in care transitions, including to certify that the beneficiary's life expectancy is six months or less before enrollment (42 CFR 418.22-26). In 2019, 52% of Medicare decedents received hospice care and Medicare spent $20.9 billion on hospice (MedPAC 2022).

[2]Medicare Benefit Policy Manual Chapter 9.

Hospice programs are organizations of physicians, nurses, and home health aides (HHAs). According to MedPAC (2022, 2010), there were 5,058 Medicare-certified programs in 2020. Most were freestanding or based in a home health agency (91%), and the remainder were hospital-based (8%) or skilled nursing facility-based (<1%). The number of programs has more than doubled since 2001, due largely to an increase in the number of for-profit hospice programs from 765 in 2001 to 3,680 in 2020. Stevenson et al. (2015) documented that many for-profit hospice programs belong to national chains such as Vitas Healthcare, Gentiva Health Services, and Heartland Hospice, and the five largest chains served a combined 15% of patients. The number of private equity-owned hospice programs has also grown from 106 in 2011 to 409 in 2019 (Braun et al. 2021).

Hospice programs vary in size and scope. In the sample of Medicare claims we describe below, we find that hospice programs' annual patient volumes ranged from 28 patients (10th percentile), to 166 patients (50th percentile), or to 830 patients or more (90th percentile) in fiscal year 2018. The largest hospice program in the country was Vitas Healthcare Corporation of Florida, which treated nearly 30,000 patients that year. And while virtually all hospice care occurs in a patient's residence, some hospice programs operate or contract with inpatient care facilities to provide short-term respite care or acute symptom management. We find that approximately 79% of hospice programs provided at least one day of inpatient hospice care in fiscal year 2018. Likewise, some hospice programs have affiliations with nursing homes to treat some patients in those settings (Stevenson et al. 2018). Patient characteristics—including health status—can vary considerably between programs (e.g., Rosenkranz et al. 2024; Furuno et al. 2020; Dalton and Bradford 2019; Gandhi 2012).

Care quality is an important determinant of patient welfare in hospice. But measurement is complicated by the fact that hospice enrollees' health can deteriorate quickly after hospice enrollment: the median length-of-stay in hospice care was just 18 days in 2020 (MedPAC 2022). Existing evidence on hospice quality often relies on survey-based measures, such as surviving family members' perceptions (e.g., Harrison et al. 2022; Anhang Price et al. 2015; Teno et al. 2011). In 2022, the Centers for Medicare and Medicaid Services (CMS) began reporting the Hospice Care Index (HCI), a composite measure of hospice quality based hospice claims data, including records of skilled nursing visits and burdensome transitions, which are defined as live discharges from hospice care followed by hospital stays (e.g., CMS 2024; Plotzke et al. 2022). Hospice care is associated with higher quality of care, fewer unmet needs, and fewer hospitalizations at end-of-life

(e.g., Gruber et al. 2023; Zuckerman et al. 2015; Kelley, et al. 2013; Teno et al. 2011).

2.2 Reimbursement rates and the cap

Medicare pays hospice programs a flat rate for each patient-day. The rate is determined by a base payment and a wage index adjustment. The base payment primarily depends on the patient's level of hospice care.[3] There are four levels of hospice care: routine home care (RHC), continuous home care (CHC), inpatient respite care (IRC), and general inpatient care (GIC).[4] Base payments for GIC and CHC are approximately 4-6x higher than base payments for RHC and IRC. However, virtually all patient-days are associated with RHC (98%). Since 2016, base payments have also been higher for each patient's first 60 days in hospice care. The base payment for RHC varied from $102 in 2001 to $193 (during the first 60 days) and $151 (after the first 60 days) in 2018.

The base payment is adjusted to account for geographic variation in health care labor costs. In particular, a portion of the base payment is multiplied by a so-called "wage index" (WI) that varies between core-based statistical areas (CBSAs) and is updated annually. On average in 2018, programs operating below the 10th percentile of the WI earned $141 per patient-day and programs operating above the 90th percentile of the WI earned $206 per patient-day.

After paying hospice programs a flat rate per patient-day, Medicare limits their average annual revenue with a so-called "aggregate payment cap." In 2023, the cap was $32,487, meaning that a hospice program that served 100 patients that year could have received no more than $3,248,700 from Medicare.[5] Any excess revenue is a cap liability that must be repaid at the end of each fiscal year.[6] In practice, audits conducted by the Office of Inspector General (OIG) suggest that CMS recovers 73-80% of each year's cap liabilities within a few years. But these audits and some press reports also suggest that cap liabilities linked to closed programs are unlikely to be collected (OIG 2022a; OIG 2021; Waldman 2012). Hospice programs are generally aware of the cap and can

[3]Since 2014, CMS has reduced base payment rates by 2% for hospice programs that have not complied with certain quality reporting requirements (Federal Register, vol. 80(151): 47207).

[4]RHC is intended for patients who are "generally stable;" IRC is "temporary care provided [...] so that the patient's caregiver can take some time off;" and both CHC and GIC are "crisis-like level[s] of care for short-term management of out of control patient pain[.]" See "Hospice levels of care" available at medicare.gov.

[5]The cap is approximately 180 days times the base daily payment rate for RHC. See Federal Register, vol. 48(163): 38156-7, Federal Register, vol. 80(151): 47147, and CMS Transmittal 11542, dated August 4, 2022.

[6]It is treated as a debt owed to the federal government and accrues interest. It may be forwarded to the Treasury Department for debt collection. See 42 CFR 418.308 and the Medicare Financial Management Manual, Chapter 4.

predict a large portion of their cap liabilities *ex interim* using their billing records.

The cap may be rationalized by concerns that provider-induced demand or moral hazard would otherwise lead to excessive hospice care. The provider-induced demand theory suggests that financially motivated health care providers may leverage informational asymmetries between themselves and their patients to persuade their patients to accept unnecessary health care services. Their scope for providing inefficient care is greatest in clinical "gray areas," where the expected harm to overutilization is low and uncertainty is high. For instance, studies have shown that the quantity of medical imaging is particularly sensitive to providers' reimbursement rates (Clemens and Gottlieb 2014; Lee and Levy 2012). Moral hazard can magnify the effects of provider-induced demand when patients are insulated from the marginal cost of care by health insurance.

Hospice care may be one such gray area. Death is unpredictable among Medicare beneficiaries (e.g., Einav et al. 2018), it is difficult for health care providers to determine when hospice care is needed most (Sack 2007), and there is virtually no patient cost-sharing under the Medicare hospice benefit.[7] There may therefore be scope for financially motivated hospice programs to raise revenue by persuading Medicare beneficiaries on the margin to enroll too early or remain enrolled too long after their condition improves. Recent press reports suggest this sometimes occurs in practice. For instance, Rao (2011) reported that an administrator for a Texas-based hospice chain "strongly encouraged employees to find a way to keep patients as long as possible" (internal quotations omitted) and whistleblowers have made similar allegations elsewhere (Kofman 2022; Rao 2011). Under these circumstances, the cap can in theory cut costs by reducing payments to hospice programs that provide long stays. But the cap's effectiveness depends on how the programs respond.

3 Cap-induced financial incentives to churn patients

The following stylized model motivates our analysis. We discuss some omitted details in appendix A and account for them in our empirical work. Define the indices j for programs and t for fiscal years. Let Cap_t be year t's cap. For each program-year, let Payments_{jt} be the program's gross Medicare revenue, let Census_{jt} be its patient census, and let $\text{PPP}_{jt} := \text{Payments}_{jt}/\text{Census}_{jt}$ be its

[7] Patients pay up to $5 per prescription and 5% of the IRC payment rate. See the Medicare Hospice Benefit Manual.

average annual revenue ("payments per patient"). Define the cap liability:

$$\text{Liability}_{jt} := \max\left\{\left(\underbrace{\frac{\text{Payments}_{jt}}{\text{Census}_{jt}}}_{\text{PPP}_{jt}} - \text{Cap}_t\right) \cdot \text{Census}_{jt}, 0\right\} = \max\left\{\text{Payments}_{jt} - \underbrace{\text{Cap}_t \text{Census}_{jt}}_{\substack{\text{MRR} \\ \text{Given Census}}}, 0\right\} \quad (1)$$

Equation (1) shows that the cap created a MRR given a program's patient census—any gross Medicare revenue in excess of the MRR is a cap liability and must be repaid to Medicare. Our central theoretical observation is that programs can reduce their cap liabilities by churning their patients: live discharges can reduce gross Medicare revenue and enrollments can increase the MRR.

These cap-related financial incentives change sharply during the transition between fiscal years. During the last weeks of a fiscal year t, programs are more certain of their cap liability and have less time to make compensating adjustments. Consequently, the expected benefit of a compensating adjustment is higher during fiscal year t's last weeks than fiscal year $t+1$'s first weeks. Since cap liabilities are paid after each fiscal year, time discounting may magnify this change by making fiscal year t's outcome more salient than fiscal year $t+1$'s outcome at the end of fiscal year t.

In the remainder of this section, we elaborate on this observation by describing how programs' financial incentives to churn patients vary because of how Medicare measures Payments_{jt} and Census_{jt}. We also describe how a 2012 change to the formula that determines Census_{jt} may have influenced their financial incentives to churn patients.

3.1 Notation and setup

Consider a program j deciding whether to enroll or live discharge a patient i on a day d near the end of a fiscal year t. (Let d be the number of days remaining in the fiscal year.) Let $g \in \{\text{E}, \text{LD}\}$ denote whether the program is facing the enrollment or live discharge decision. Let P_{jt} be its daily payment rate. Let $A \in \{0, 1\}$ be its action (enroll vs. do not enroll or live discharge vs. do not live discharge). Let $D \in \{0, 1\}$ indicate whether it is on track to exceed the cap.

The program's choice of how to act will depend on its payoffs. Let $\tilde{\pi}^g_{ijt}(A, D)$ be its payoff function, which depends on the effect of its action on its cap liability and other factors. Define

$\tilde{\pi}^g_{ijt}(A, D) := \text{Liability}^g_{ijt}(A, D) + \varepsilon^g_{ijt}(A)$, where $\varepsilon^g_{ijt}(A)$ includes the program's gross profit for treating this patient in this fiscal year, the action's fixed cost, and the present value of payoffs related to this patient in future fiscal years. Let $\pi^g_{ijt}(D) := \tilde{\pi}^g_{ijt}(1, D) - \tilde{\pi}^g_{ijt}(0, D)$ be the action's net payoff to the program and let the cap-induced financial incentive for the action be:

$$\Delta\pi^g_{ijt} := \pi^g_{ijt}(1) - \pi^g_{ijt}(0) \tag{2}$$

We assume that the program will act in expectation over the distribution of patient characteristics.

Medicare determines the functional form of equation (2). It varies with the fiscal year because of a 2012 change to how Medicare measures patient censuses. Before 2011, Medicare used the so-called "streamlined" method. Under this method, patients who are only ever enrolled with one program contribute 1 to their program's census in the fiscal year that they enroll for the first time. They contribute 0 to their program's census in all subsequent fiscal years. However, patients who are ever enrolled with more than one program contribute to each program-year's census the fraction of their LLOS spent at that program in that fiscal year. Since 2012, Medicare has used the so-called "proportional" method at most programs.[8] Under this method, all patients contribute to each program-year's census the fraction of their LLOS spent at that program in that fiscal year. Programs make a preliminary estimate of their cap liabilities at the end of each fiscal year and adjust it upward for up to three years afterward as their enrollees continue using hospice care.

For brevity, we assume that a marginal enrollment or live discharge would not create a cap liability for a program not on track to exceed the cap. Therefore $\text{Liability}^g_{ijt}(A, 0) = 0$. We also assume that a marginal enrollment or live discharge would not eliminate the cap liability of a program on track to exceed the cap. Programs very near the cap have qualitatively similar—but correspondingly milder—financial incentives to those we describe here. Finally, we assume that the level of hospice care is fixed because RHC is provided for 98% of an average hospice program's patient-days. We hypothesize that programs on track to exceed the cap may reduce GIC and CHC rates because they are more resource intensive but provide no marginal net revenue because of the cap.

[8]Most programs were automatically switched to the proportional method in 2012. However, a small and decreasing number of programs continued using the streamlined method (only 486 programs in 2013). We do not observe which programs used which methods, so we assume that all programs switched to the proportional method in 2012.

3.2 Cap-induced financial incentives to enroll new patients

We begin with the program deciding whether to enroll a new patient. For brevity, we assume that the patient had never previously been enrolled in hospice care. Let LOS_{ijt} be the patient's length-of-stay at hospice j during fiscal year t, let LLOS_i be the patient's LLOS, and let $W_i \in \{0, 1\}$ indicate whether the patient has ever or will ever be enrolled with a different program. Note that $\text{Liability}^{\text{E}}_{ijt}(0, D) = 0$ because the patient's effect on the program's cap liability is zero if the program does not enroll the patient. Under these assumptions:

$$\Delta\pi^{\text{E}}_{ijt} = -\text{Liability}^{\text{E}}_{ijt}(1,1) = \begin{cases} \text{Cap}_t - P_{jt}\text{LOS}_{ijt} & \text{if } t \leq 2011 \text{ and } W_i = 0 \quad (3a) \\ \dfrac{\text{LOS}_{ijt}}{\text{LLOS}_i}\text{Cap}_t - P_{jt}\text{LOS}_{ijt} & \text{if } t \leq 2011 \text{ and } W_i = 1 \quad (3b) \\ \dfrac{\text{LOS}_{ijt}}{\text{LLOS}_i}\text{Cap}_t - P_{jt}\text{LOS}_{ijt} & \text{if } t \geq 2012 \quad (3c) \end{cases}$$

We make three observations about $\Delta\pi^{\text{E}}_{ijt}$ near the end of the fiscal year. First, the cap-induced financial incentive to enroll a new patient varies with the patient's characteristics, but would be positive for most patients. As we discuss below, the average hospice earned \$146 per patient-day and faced a cap of \$24,047 during our sample period. Therefore, equation (3a) is positive near the end of the fiscal year and equations (3b) and (3c) are positive as long as LLOS_i is not too large.[9] In our empirical analysis, we will examine whether programs on track to exceed the cap enroll more patients than otherwise similar programs not on track to exceed the cap contemporaneously. We will also examine the characteristics of the marginal enrollees.

Second, the cap-induced financial incentive to enroll a new patient was larger before 2011 because $\text{Cap}_t \geq \frac{\text{LOS}_{ijt}}{\text{LLOS}_i}\text{Cap}_t$. In our empirical analysis, we will examine whether programs on track to exceed the cap enrolled more patients before 2011 than they have since 2012, relative to otherwise similar programs not on track to exceed the cap contemporaneously.

Third, the cap-induced financial incentive to enroll a new patient varies with the day d of the fiscal year. Before 2011, $\Delta\pi^{\text{E}}_{ijt}$ was generally decreasing in LOS_{ijt} because $W_i = 0$ for most patients. This suggests that $\Delta\pi^{\text{E}}_{ijt}$ generally increased as $d \to 0$ because $\text{LOS}_{ijt} \leq d$. Since 2012, holding

[9] In practice, the programs's cap liability payment will depend on the patient's LLOS through the end of the fiscal year, not LLOS_i itself. It will subsequently make additional payments to account for updates to its patients' lifetime hospice utilization. Therefore, time discounting will diminish the impact of a large LLOS_i on present payoffs.

LLOS_i fixed, the cap-induced financial incentive to enroll a new patient has converged to zero as $d \to 0$. Figure 1 illustrates how $\Delta\pi^{\text{E}}_{ijt}$ varies with LLOS_i and d for typical values of P_{jt} and Cap_t. In our empirical analysis, we will examine enrollment trends at programs on track to exceed the cap near the end of a fiscal year.

3.3 Cap-induced financial incentives to live discharge patients

We continue with the program deciding whether to live discharge a patient. For brevity, we assume that the patient's first hospice enrollment occurred at this program during this fiscal year. The program's cap-induced financial incentive to live discharge the patient depends on the effect of the live discharge on the patient's lifetime utilization of hospice care. Let $\text{LOS}_{ijt}(A)$ be the patient's length-of-stay at hospice j during fiscal year t, let $\text{LLOS}_i(A)$ be the patient's LLOS, and let $W_i(A) \in \{0, 1\}$ indicate whether the patient has or will ever be enrolled with another program. For each of these terms K, define $\Delta K := K(1) - K(0)$ and $\%\Delta K := \Delta K / K(0)$. For brevity, we assume that $\Delta W_i \geq 0$. Under these assumptions:

$$\Delta\pi^{\text{LD}}_{ijt} = -\left(\text{Liability}^{\text{LD}}_{ijt}(1,1) - \text{Liability}^{\text{LD}}_{ijt}(1,0)\right) =$$

$$\begin{cases} -\Delta\text{LOS}_{ijt} P_{jt} & \text{if } t \leq 2011 \text{ and } W_i(1) = 0 \text{ and } W_i(0) = 0 \quad (4a) \\ \left(\dfrac{\text{LOS}_{ijt}(1)}{\text{LLOS}_i(1)} - 1\right)\text{Cap}_t - \Delta\text{LOS}_{ijt} P_{jt} & \text{if } t \leq 2011 \text{ and } W_i(1) = 1 \text{ and } W_i(0) = 0 \quad (4b) \\ \left(\dfrac{\text{LOS}_{ijt}(1)}{\text{LLOS}_i(1)} - \dfrac{\text{LOS}_{ijt}(0)}{\text{LLOS}_i(0)}\right)\text{Cap}_t - \Delta\text{LOS}_{ijt} P_{jt} & \text{if } t \leq 2011 \text{ and } W_i(1) = 1 \text{ and } W_i(0) = 1 \quad (4c) \\ \left(\dfrac{\text{LOS}_{ijt}(1)}{\text{LLOS}_i(1)} - \dfrac{\text{LOS}_{ijt}(0)}{\text{LLOS}_i(0)}\right)\text{Cap}_t - \Delta\text{LOS}_{ijt} P_{jt} & \text{if } t \geq 2012 \quad (4d) \end{cases}$$

In equation (4b), note that $\left(\frac{\text{LOS}_{ijt}(1)}{\text{LLOS}_i(1)} - 1\right) \leq 0$. In equations (4c) and (4d), note that $\left(\frac{\text{LOS}_{ijt}(1)}{\text{LLOS}_i(1)} - \frac{\text{LOS}_{ijt}(0)}{\text{LLOS}_i(0)}\right) \geq 0$ if and only if $\%\Delta\text{LOS}_{ijt} \geq \%\Delta\text{LLOS}_i$.

We make three observations about $\Delta\pi^{\text{LD}}_{ijt}$ near the end of the fiscal year. First, the cap-induced financial incentive to live discharge a patient may be positive or negative and depends on the patient's characteristics. If $\Delta\text{LOS}_{ijt} \leq 0$—which is plausible for most patients—then equation (4a) is positive. But equations (4b), (4c), and (4d) may be positive or negative in ways that depend on unknown potential outcomes. We hypothesize that the program can predict to some extent which patients would reduce its cap liability if they were live discharged. For instance, before 2011, it

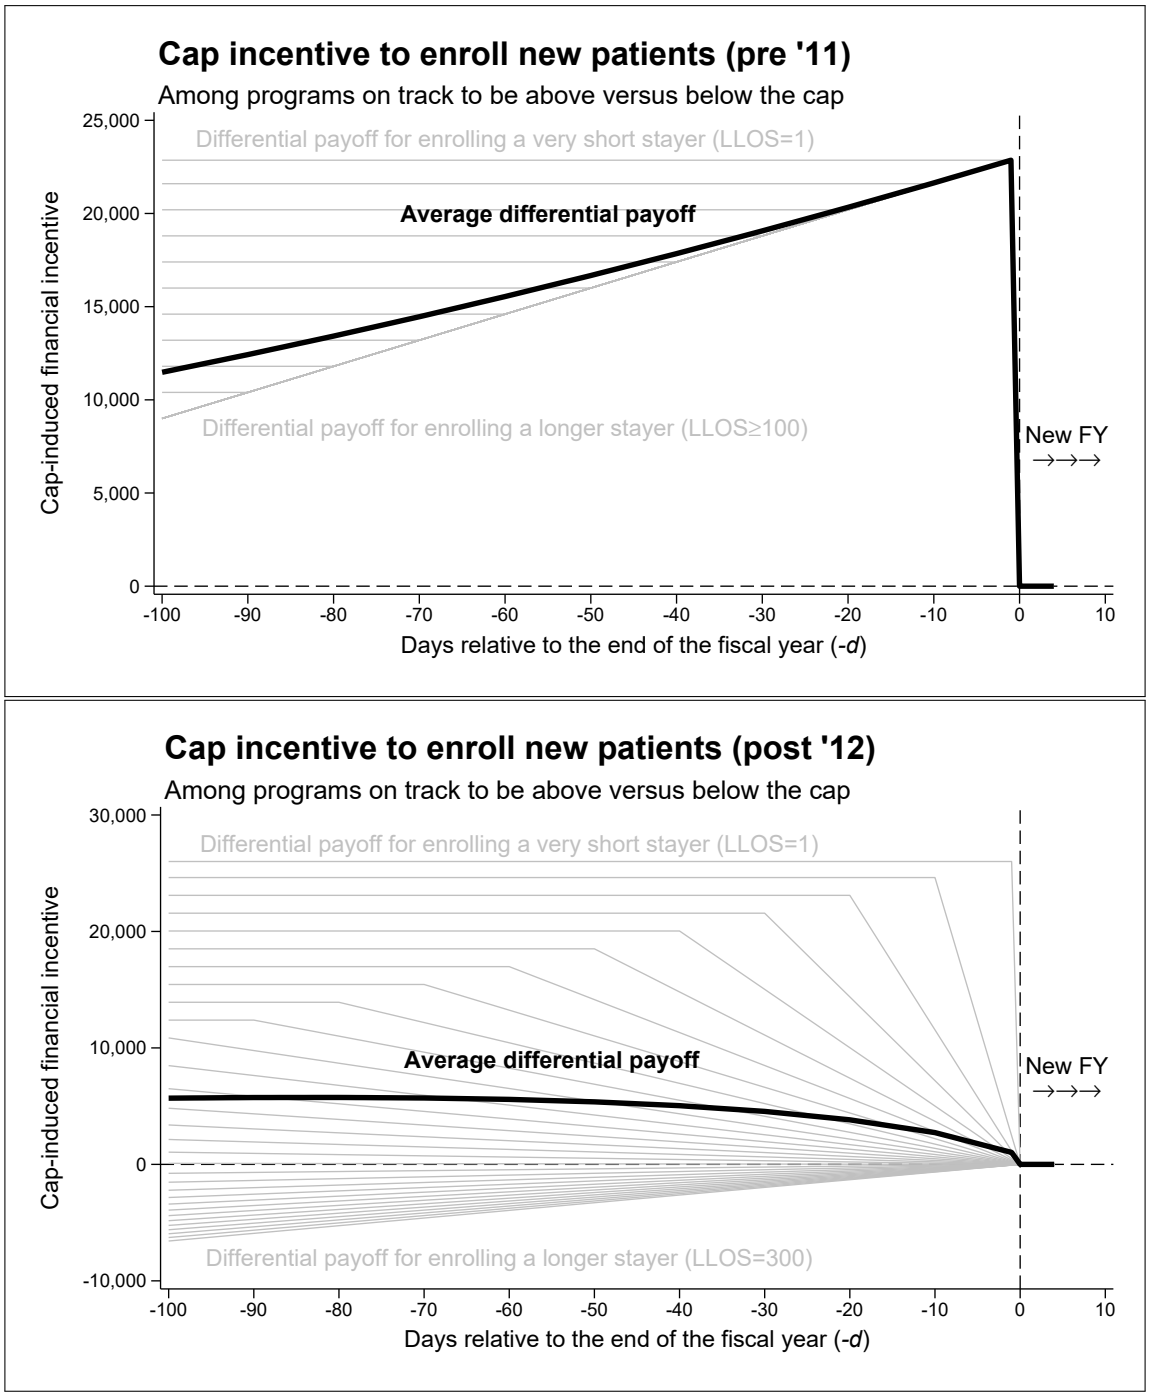


Fig. 1. These figures illustrate the cap-induced financial incentives to enroll new patients near the end of a fiscal year. The top panel plots the incentive for fiscal years prior to 2011 and patients with $W_i = 0$ (corresponding to equation (3a)). The bottom panel plots the incentive for fiscal years after 2012 or patients with $W_i = 1$ (corresponding to equations (3b) and (3c)). Both panels show that programs on track to exceed the cap in an outgoing fiscal year generally benefit from enrolling new patients relative to otherwise similar hospice programs not on track to exceed the cap contemporaneously.

may identify patients who are not likely to re-enroll elsewhere if they were live discharged. After 2012, it may identify patients who are likely to remain enrolled much longer but-for a live discharge, such as those for whom $\text{LOS}_{ijt}(1) \approx \text{LLOS}_i(1)$ but $\text{LOS}_{ijt}(0)/\text{LLOS}_i(0) << 1$. In our empirical analysis, we will examine whether programs on track to exceed the cap live discharge more patients than otherwise similar programs not on track to exceed the cap contemporaneously. We will also examine the characteristics of marginally live discharged patients.

Second, the 2012 policy change increased the cap-induced financial incentive to live discharge some patients, but decreased it for others. Since only 5% of patients—including only 30% of live discharged patients—ever enroll with one hospice, we hypothesize that the number of patients for whom equation (4a) is positive is larger than the number of patients for whom equation (4d) is positive. In our empirical analysis, we will examine whether programs on track to exceed the cap live discharged more patients before 2011 than they have since 2012, relative to otherwise similar programs not on track to exceed the cap contemporaneously.

Third, the cap-induced financial incentive to live discharge a patient varies with the day d of the fiscal year. But while $\Delta\text{LOS}_{ijt} \to 0$ as $d \to 0$, the first terms in equations (4b), (4c), and (4d) converge to functions of unknown potential outcomes. In our empirical analysis, we will examine live discharge trends at programs on track to exceed the cap near the end of a fiscal year.

4 Data

Our primary data are the 100% Medicare hospice claims spanning 2000-2019. We extracted patient and program identification numbers and claim start and end dates associated with each claim. We used these data to construct various patient and program panel datasets.[10] Our primary analytic samples are (program, day)- and (program, fiscal year)-level datasets spanning 2001-2018.[11]

We used the claims to measure hospice programs' annual gross Medicare payments and patient censuses. Payments are observed directly in the claims data.[12] Patient censuses are determined

[10]Almost all (patient, day) observations (99.98%) are associated with at most one program. Patients may be associated with multiple programs on days that they switch between programs. Such patient-days count toward both programs' cap censuses. See Medicare hospice transmittal 156.

[11]We use 2000 and 2019 data to reduce left- and right-censoring in measures of patients' lifetime hospice utilization.

[12]For claims that spanned several days, we distributed payments evenly across the days associated with that claim. We accounted for several details in the payment formula, including the effect of a 2013 sequestration and how days are assigned to fiscal years for the purposes of counting annual payments.

by each patient's lifetime utilization of hospice care and the counting method. We calculated each patient-day's contribution to its program-year's patient census using the streamlined method before 2011 and the proportional method thereafter. In both cases, we used records of each patient's hospice utilization during the following three fiscal years because a program-year's cap liability may be revised for up to three additional fiscal years as its patients continue using hospice care.[13]

We combined our measures of annual Medicare payments and patient censuses with publicly available data published in the Federal Register or CMS's transmittals to hospice programs about each fiscal year's cap amount and WI.[14] We used these data to calculate each program-year's average annual revenue and cap liability. We also computed for each program-year-day the program's running average annual revenue and cap liability on that day of the fiscal year.

We also relied on several additional data sources. First, we used the hospice claims and 2000-2019 Master Beneficiary Summary Files (MBSF) to measure aspects of each patient's lifetime hospice utilization. We defined each patient's first enrollment date as the first day of their first observed claim. We defined a live discharge as any transition between distinct programs or any break in a patient's hospice claims spanning 1+ days that does not coincide with their date of death as reported in the MBSF. We defined each patient's LLOS as their total number of days in hospice care during 2000-2019. We defined their lifetime number of enrollments with distinct programs analogously. We determined each patient-day's level of hospice care (RHC, CHC, GIC, or IRC) and whether a staff visit occurred using the revenue center codes associated with each claim.

Second, we used the hospice claims and MBSF to identify several other patient characteristics. We identified each hospice patient's date of birth, race, sex, and date of death in the MBSF.[15] We defined a decedent to be a patient whose date of death was on or before the end of our sample period, December 31, 2019. We defined each patient's remaining days of life (RDOL) as the difference between their enrollment date and their earlier of their date of death or the end of our sample period. We defined each patient as dying during hospice care if a hospice claim encompasses their date of death. We also used the MBSF to identify each patient-year's county of residence, which we used

[13]Among patients whose first observed patient-day was between FY 2004 and FY 2015, 92% of their patient-days are associated with either their first observed patient-day's fiscal year or the following fiscal year. We accounted for several details in the patient census formulas, including how patients were counted during the short transition between counting methods in 2011 and how days are assigned to fiscal years for the purposes of counting patients.

[14]We accounted for the fact that the cap assigned to new hospice programs is different during their first fiscal year.

[15]It is not common for one patient to have multiple distinct realizations of these variables. When that happens, we assigned them to whichever realization was reported first.

to measure geographic overlap between each program's patients and patients at programs with cap liabilities. We refer to this as the program's proximity to over-cap programs.

Third, we used the 2000-2019 provider-of-service (POS) files to identify hospice program characteristics. In particular, we identified each program's name and state, and whether the program was for-profit, not for-profit, or government-owned. We also used the programs' original participation dates and their earliest observed hospice claims to determine their starting years and ages. We defined a hospice program to be exiting in a given fiscal year if we do not observe that program's Medicare provider certification number (CCN) in the following year's claims.[16]

Fourth, we used the 2000-2019 Medicare Provider Analysis and Review files (MedPAR) and the Minimum Dataset (MDS). MedPAR contains records of Medicare beneficiaries' hospital and skilled nursing facility (SNF) stays.[17] We extracted the admission and discharge date associated with each MedPAR claim to determine patient-days spent in a hospital or SNF. The MDS contains health assessments conducted by NHs when a patient begins and ends a NH stay, and at regular intervals during their NH stay. We identified patient-days spent at a NH by linking sequences of MDS assessments and merging them with MedPAR's records of SNF stays.

See appendix B for more information about these data.

5 Descriptive statistics

The following sections summarize our data as they relate to the cap. We report other patient and provider descriptive statistics in appendix C.

5.1 The distribution of cap liabilities

Table 1 summarizes the programs' patient volumes, Medicare payments, and cap liabilities. Column (1) shows that hospice programs exceeded the cap in 11% of program-years during 2001-2018. Column (2) shows that 2,214 programs (39%) exceeded the cap at least once and their cap liabilities were $490,000 per year on average. Columns (2) and (6) show that programs that exceeded the cap

[16] A Medicare provider agreement is terminated when the corresponding provider exits the Medicare market or experiences a change of ownership to an owner who does not accept the agreement. To the best of our knowledge, there is no taxonomy of hospice program CCN terminations, so we cannot distinguish exits from other events.

[17] We linked MedPAR to the POS files to determine which providers were hospitals and SNFs.

treated fewer unique patients than other programs, but for more patient-days on average. Columns (3) and (4) show that programs just above the cap also had significant cap liabilities—$163,000 per year, on average. MedPAC (2020) proposed reducing the cap by 20%. Column (5) reports descriptive statistics for programs with average annual revenue below but within 20% of the cap. It shows that they treated more than twice as many patients as programs that had exceeded the cap.

Figure 2 plots the distribution of $\text{PPP}_{jt}/\text{Cap}_t$ and table A1 reports the McCrary test of the hypothesis that the distribution of $\text{PPP}_{jt}/\text{Cap}_t$ is continuous at 1 (McCrary 2008). There is no evidence of bunching around 1, indicating that there is no detectable mass of programs that perfectly adjust to the cap.[18] This implies that programs' marginal adjustment costs are high because even programs with average annual revenue between 100-110% of the cap had cap liabilities of $163,000 per year on average. A notable exception is Vitas Healthcare Corporation of Florida, which we find in appendix C had an average annual revenue within 4.4% of the cap on average during 2001-2018. A recent OIG report suggests that it may have inappropriately billed CHC and GIC days during this time; if true, this suggests that the cap may have partially limited the alleged fraud (OIG 2022b).

5.2 Cap liabilities and the wage index

Next, we examined the association between the WI and cap liabilities. The WI creates geographic variation in programs' daily payment rates, but not in the cap. MedPAC (2020) proposed WI-adjusting the cap to make it "more equitable across providers." We estimated models of the form $Y_{jt} = \beta \text{WI}_{jt} + \Gamma X_{jt} + \varepsilon_{jt}$, where Y_{jt} is an outcome for program j in fiscal year t, WI_{jt} is the natural log of the WI, and X_{jt} is a vector of other characteristics, including program and year fixed effects.

Table A2 reports our results. We find that a 1% higher WI is associated with an 0.11 percentage point higher likelihood of exceeding the cap and a 4.6% higher cap liability on average. This higher cap liability is explained by higher gross Medicare revenue: we find that a 1% higher WI is associated with a 0.33% increase in gross Medicare payments per patient-day. On aggregate, we find that a 1% higher WI is associated with higher gross Medicare payments and lower patient censuses—explaining the higher cap liabilities—though these associations are not individually sta-

[18] In parallel work, Gruber et al. (2023) draw a similar conclusion using a simplified measure of cap liabilities. Table A1 reports similar results for sub-samples of for-profit programs, large programs, and programs that exceeded the cap at least once, respectively. Although the test statistic for the sub-sample of programs that exceeded the cap at least once is statistically significant, it is not of the expected sign.

		Above-cap			Below-cap	
	All	All	W/I 10%		≥80%	All
	(1)	(2)	(3)	(4)	(5)	(6)
Patient volume and cap census						
Patient-days	23,493	20,774	26,495	36,372	33,864	23,846
Unique patients	353	178	250	375	360	376
Patient-days per unique patient	71	122	107	101	97	65
Cap census	270	105	157	248	241	292
Cap census contr. per unique pat.	0.73	0.55	0.60	0.63	0.64	0.75
Frac. pat.-days stl. (pre-'11)	0.85	0.66	0.71	0.74	0.76	0.87
Frac. pat.-days prop. (pre-'11)	0.15	0.34	0.29	0.26	0.24	0.13
Gross Medicare payments						
Payments ($K)	3,640	3,023	4,028	5,819	5,288	3,720
Payments per pat.-day ($)	146	148	150	149	149	146
Payments per unique pat. ($)	10,306	17,648	15,771	14,888	14,260	9,353
PPP ($)	15,241	33,598	26,396	23,828	22,354	12,856
Cap liabilities and Medicare payments net of cap liabilities						
Cap ($)	24,047	24,932	25,217	25,137	25,093	23,932
PPP/Cap	0.63	1.35	1.05	0.95	0.89	0.53
Cap liability ($K)	56	490	163	0	0	0
Net payments ($K)	3,583	2,532	3,865	5,819	5,288	3,720
Net payments per pat.-day ($)	143	118	143	149	149	146
Place in the distribution of cap liabilities						
1[Above cap]	0.11	1.00	1.00	0.00	0.00	0.00
1[100-110% of cap]	0.04	0.31	1.00	0.00	0.00	0.00
1[90-100% of cap]	0.05	0.00	0.00	1.00	0.43	0.06
1[80-100% of cap]	0.12	0.00	0.00	1.00	1.00	0.13
1[Below cap]	0.89	0.00	0.00	1.00	1.00	1.00
N (program-years)	57,412	6,602	2,018	2,926	6,824	50,810
Unique programs	5,685	2,214	1,279	1,605	2,517	5,300

Tab. 1. This table reports cap-related descriptive statistics about the (program, fiscal year)-level data. On average, programs that exceed the cap have substantial cap liabilities, treat fewer patients, and earn 20% less per patient-day net of their cap liabilities ($118 vs. $148). Their patients contribute less to their annual cap censuses and have longer LOS in hospice care.

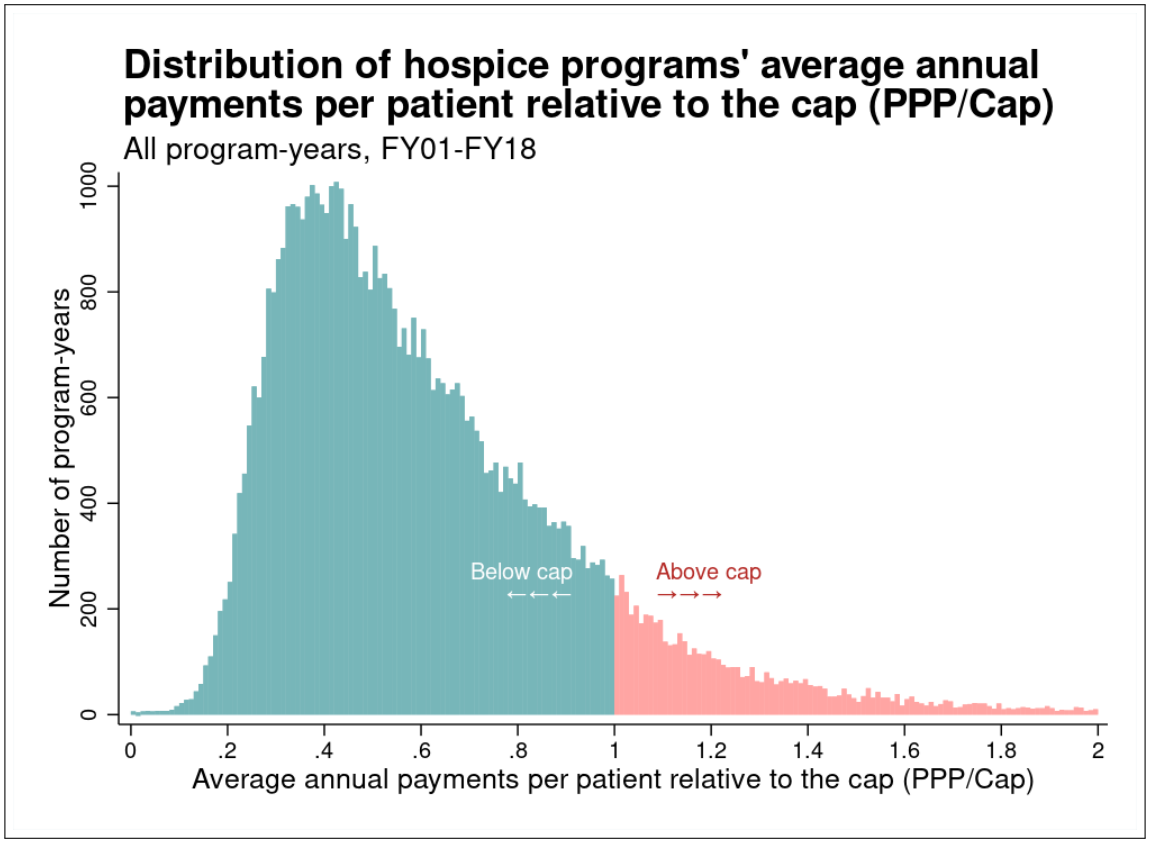


Fig. 2. This figure plots the distribution of $\text{PPP}_{jt}/\text{Cap}_{jt}$—i.e., hospice programs' average annual payments per patient relative to the cap. The figure is censored at $\text{PPP}_{jt}/\text{Cap}_{jt} = 2$ for visual clarity. There is no visual indication that hospice programs' average annual payments per patient bunch around the cap. We report McCrary test results in table A1.

tistically significantly different from zero. In sum, these estimates support MedPAC's claim that WI-adjusting the cap would improve equity across programs because a portion of the WI adjustment intended to compensate programs for higher local labor costs is repaid in cap liabilities.

5.3 What happens after programs exceed the cap?

Next, we examined how programs change after exceeding the cap. MedPAC (2020) proposed reducing the cap by 20%. Figure 2 suggests that in the absence of a compensating response, this proposal would cause a significant number of programs to exceed the cap more often. What happens to programs that exceed the cap in future fiscal years? We estimated models of the form $Y_{jt} = \beta 1[\text{Over Cap}]_{jt-1} + \Gamma X_{jt} + \varepsilon_{jt}$, where Y_{jt} is an outcome for program j in fiscal year t, $1[\text{Over Cap}]_{jt-1}$ indicates that the program exceed the cap in the last fiscal year, and X_{jt} is a vector

of other characteristics, including program and year fixed effects.

Table A3 reports our results. It shows that exceeding the cap in one fiscal year is associated with a 19.3 percentage point higher likelihood of exceeding it again in the following fiscal year, and a 246% higher average cap liability. It also shows that programs that exceed the cap in one fiscal year have 10.1% lower patient censuses but 10.0% more patient-days in the following fiscal year, on average. Importantly, it shows that CCNs associated with programs that exceed the cap are 5.6 percentage points more likely to be terminated in the next fiscal year on average. By comparison, the baseline termination rate is 2.5% per year on average. We caution that causal inference is complicated by the possibility that unobserved factors—such as poor management or decreasing residual demand for hospice care—may confound our estimates. But together with our previous finding that programs below and within 20% of the cap treat more patients on average, these associations raise concerns about the potential effect of reducing the cap on market structure and hospice access.

6 Do programs churn patients because of the cap?

6.1 Research design

6.1.1 Setup

The foregoing section shows that programs exceed the cap by wide margins and there is no bunching in the distribution of average annual revenue. Together, these findings suggest that marginal adjustment costs are high by the end of the fiscal year. But do programs end a fiscal year with cap liabilities without making adjustments? Or do their adjustments fall short of eliminating their cap liabilities? In this section, we investigate how programs on track to exceed the cap in an outgoing fiscal year adjust to the cap by leveraging variation in cap-related financial incentives generated by the policy's nonlinear design and the transition between fiscal years.

In particular, we treat the transitions between fiscal years 2001-2002, ..., 2018-2019 as eighteen distinct events and examine outcome trends during the 180-days surrounding each event ("event windows"). For most outcomes, the transition date is the start of the fiscal year, which was November 1 until 2017 and October 1 thereafter. For others, the appropriate transition date differs because of Medicare's census counting methods. Under the streamlined method (before 2011), patients who

were never treated by more than one program were counted until September 27. Other patients were counted until the end of the fiscal year. Under the proportional method (after 2012), all patients were counted until the end of the fiscal year. For census related outcomes, the transition date is September 28 until 2011 (since most patients are treated by one program at most) and the start of the fiscal year thereafter. We call it the transition between "primary census counting periods."

We define a program to be on track to exceed the cap in an outgoing fiscal year if its running annual payments exceed its running patient census at the start of the event window. It is plausible that by this time, programs can accurately forecast an impending cap liability.[19] We use a DID research design to compare outcomes between programs on track to exceed the cap in an outgoing fiscal year and observably similar programs not on track to exceed the cap contemporaneously.

6.1.2 Estimation

Define the indexes e for events, j for programs, and t for fiscal years. Let t_e be event e's outgoing fiscal year and let Cap_e be the cap in t_e. Let PPP_{ej} be program j's running average annual revenue at the start of e's event window. Let the set of programs on track to exceed the cap in t_e be $\mathcal{T}_e := \{j : \text{PPP}_{ej} \geq \text{Cap}_e\}$ and let $\mathcal{U}_e$ be the others. Let $W := \{..., -2, -1, 0, 1, 2, ...\}$ enumerate 7-day intervals relative to the event's transition date and let w index elements of W. Finally, define the static treatment indicator $D^{\text{Sta.}}_{ejw} := 1[j \in \mathcal{T}_e]1[w < 0]$ and the dynamic treatment indicator $D^{\text{Dyn.}}_{ejww'} := 1[j \in \mathcal{T}_e]1[w = w']$. We estimate stacked DID regression models of the form:

$$Y_{ejw} = \beta D^{\text{Sta.}}_{ejw} + \text{FE}_{ej} + \text{FE}_{ew} + \varepsilon_{ejw} \tag{5}$$

$$Y_{ejw} = \sum_{w' \in W \setminus \{0\}} D^{\text{Dyn.}}_{ejww'} \beta_{w'} + \text{FE}_{ej} + \text{FE}_{ew} + \varepsilon_{ejw} \tag{6}$$

where Y_{ejw} is an outcome, FE_{ej} is an event-program fixed effect, and FE_{ew} is an event-week fixed effect interacted with each program's state, ownership type, and WI at the start of the event window.

For binary outcomes, we estimated equations (5)- (6) using ordinary least squares (OLS) and interpret the coefficients as percentage point differences. For other non-negative outcomes, we

[19] Recall that programs cannot perfectly forecast an impending cap liability because it depends on their patients' hospice utilization for the following three fiscal years. However, we found that the correlation between each program's actual running cap liability at the start of the event window and the running cap liability it would have if cap liabilities were determined without the 3-year lookback period is 0.97. This suggests that programs could use existing billing records and a short-run projection of their patient volume to accurately forecast an impending cap liability.

estimated analogous equations using the Poisson pseudo-maximum likelihood estimator (PPML) and interpret the coefficients as percent differences (Wooldridge 1999). The stacked estimator overcomes the so-called "negative weighting" issue that can arise in two-way fixed effects (TWFE) OLS DID estimators (e.g., Goodman-Bacon 2021). The parameters β and $(\beta_w : w \in W)$ measure average trend differences during the event windows between programs in $\mathcal{T}_e$ and $\mathcal{U}_e$ in the same state, with the same ownership type, and with similar WI values (Gardner 2021; Cengiz et al. 2019).[20]

6.1.3 Identification

DID identifies the average treatment effect on the treated of the cap at the end of an outgoing fiscal year under the following assumptions. First, average outcomes among programs in $\mathcal{U}_e$ are unaffected by the cap during e's event window ("clean controls"). Second, average outcomes among programs in $\mathcal{T}_e$ are unaffected by the cap at the start of e's incoming fiscal year ("clean slate"). Third, average outcome trends between programs in $\mathcal{T}_e$ and $\mathcal{U}_e$ would have moved synchronously at the end of e's outgoing fiscal year but for the cap ("parallel trends").[21]

Figure 3 illustrates this identification strategy. It plots average enrollment rates between programs grouped at the start of the event window by their running average annual revenue relative to the cap. It shows that enrollment rates vary seasonally, and are consistently lower among programs on track to exceed the cap. These observations suggest that causal inference from within- or between-program comparisons alone may be complicated by other factors that covary with programs or time. But the figure also shows that programs with higher average annual revenue have higher differential enrollment rates at the end of an outgoing fiscal year. We attribute this variation to the cap and measure it with our regression analysis.

This approach has some limitations. First, the clean controls assumption would be violated if programs in $\mathcal{U}_e$ respond to the risk that they could ultimately exceed the cap. We therefore examine whether our estimates are sensitive to excluding programs in $\mathcal{U}_e$ with a running average annual revenue within 10% of the cap at the start of the event window. Second, the clean slate assumption

[20] Gardner (2021) shows that the stacked OLS estimator produces a weighted average of each event's trend differences, with weights that increase in the event's sample size and the parity between $|\mathcal{T}_e|$ and $|\mathcal{U}_e|$.

[21] The sets $\mathcal{T}_e$ and $\mathcal{U}_e$ contain only those (e, j) realizations such that j was observed to treat at least one patient throughout e's event window. Within each (e, j) realization, the observations $\{(e, j, w) : w \in W\}$ are balanced. That there are no compositional changes in the estimation sample within any given event implies that we may interpret $(\beta_w : w \in W)$ as the average evolution of the average treatment effect on the treated under the foregoing assumptions.

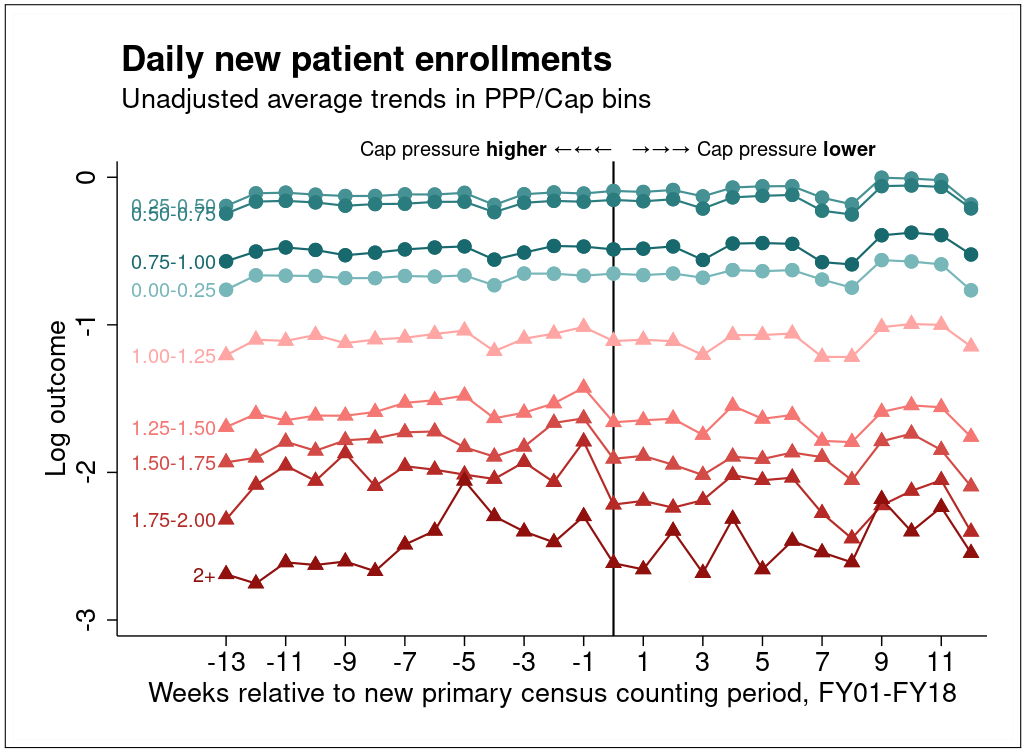


Fig. 3. This figure plots average outcome trends during the transition between fiscal years. Programs are grouped into bins based on their rolling average annual revenue just before week -13. It shows that enrollment rates vary seasonally, and are consistently lower among programs on track to exceed the cap. But it also shows that programs with higher average annual revenue have higher differential enrollment rates at the end of an outgoing fiscal year.

would be violated if programs in $\mathcal{T}_e$ enroll fewer patients in the incoming fiscal year because they responded to the cap by enrolling new patients in the outgoing fiscal year. We therefore examine whether enrollment rate trends between programs in $\mathcal{T}_e$ and $\mathcal{U}_e$ diverge during the incoming fiscal year, with enrollment rates among programs in $\mathcal{T}_e$ differentially increasing as their enrollees' pass away or exit hospice. We also investigate whether staff visit rates differentially decrease among programs in $\mathcal{T}_e$. Finally, the parallel trends assumption would be violated if programs in $\mathcal{T}_e$n operate in markets with seasonally different preferences for hospice care (e.g., due to seasonally different disease incidence rates). This concern should be mitigated by our inclusion of state, ownership type, and WI controls. We also benefit from granular time-series variation—the index w measures weeks—enabling us to examine outcome trend differences in a narrow window around the transition dates.

6.2 Results

6.2.1 Churning

Our main results document the cap's effect on churning. First, table 2 shows that average enrollment rate differences between the outgoing fiscal year's last quarter and the incoming fiscal year's first quarter are 0.056 log points (5.8%) higher among programs on track to exceed the cap in the outgoing fiscal year.[22] It also shows that average live discharge rate differences between the outgoing fiscal year's last quarter and the incoming fiscal year's first quarter are 0.042 log points (4.3%) higher among programs on track to exceed the cap in the outgoing fiscal year.

Second, we examined the impact of the 2012 reform to the census counting method. We interacted the treatment indicators D in equations (5)-(6) with an indicator equal to 0 if event e's outgoing fiscal year is 2011 or earlier, and 1 otherwise. Table 2 reports our results. It shows that programs on track to exceed the cap in an outgoing fiscal year have had lower differential enrollment and live discharge rates since 2012, consistent with the idea that the transition to the proportional counting method moderated their cap-induced financial incentives to churn patients.

Third, we examined the dynamic DID in enrollment and live discharge rates. The temporal findings are telling: figure 4 shows that enrollment rates differentially increase throughout the outgoing fiscal year's last quarter before differentially decreasing by 0.150 log points (13.9%) in the first week of the new fiscal year. Consistent with our motivating theory in section 3, this pattern is more pronounced under the streamlined method (before 2011) than it has been since. Our motivating theory also showed that live discharge dynamics depend on functions of unknown potential outcomes; and we do not observe a clear pattern in the live discharge dynamic DID estimates. Finally, we do not observe evidence against the clean slate and parallel trends assumptions.

6.2.2 Marginally enrolled and live discharged patients

We next estimated equations (5)-(6) using outcomes that describe the enrolled and live discharged patients. All observations are weighted by the number of patients who were enrolled or live discharged during the associated program-week. We report the static DID estimates in tables 3 and 4. We report the dynamic DID estimates in figures A1 and A2.

[22] Throughout the paper, we convert log differences to percent changes by: Percent = $\big(\exp(\text{Log}) - 1\big) \times 100$.

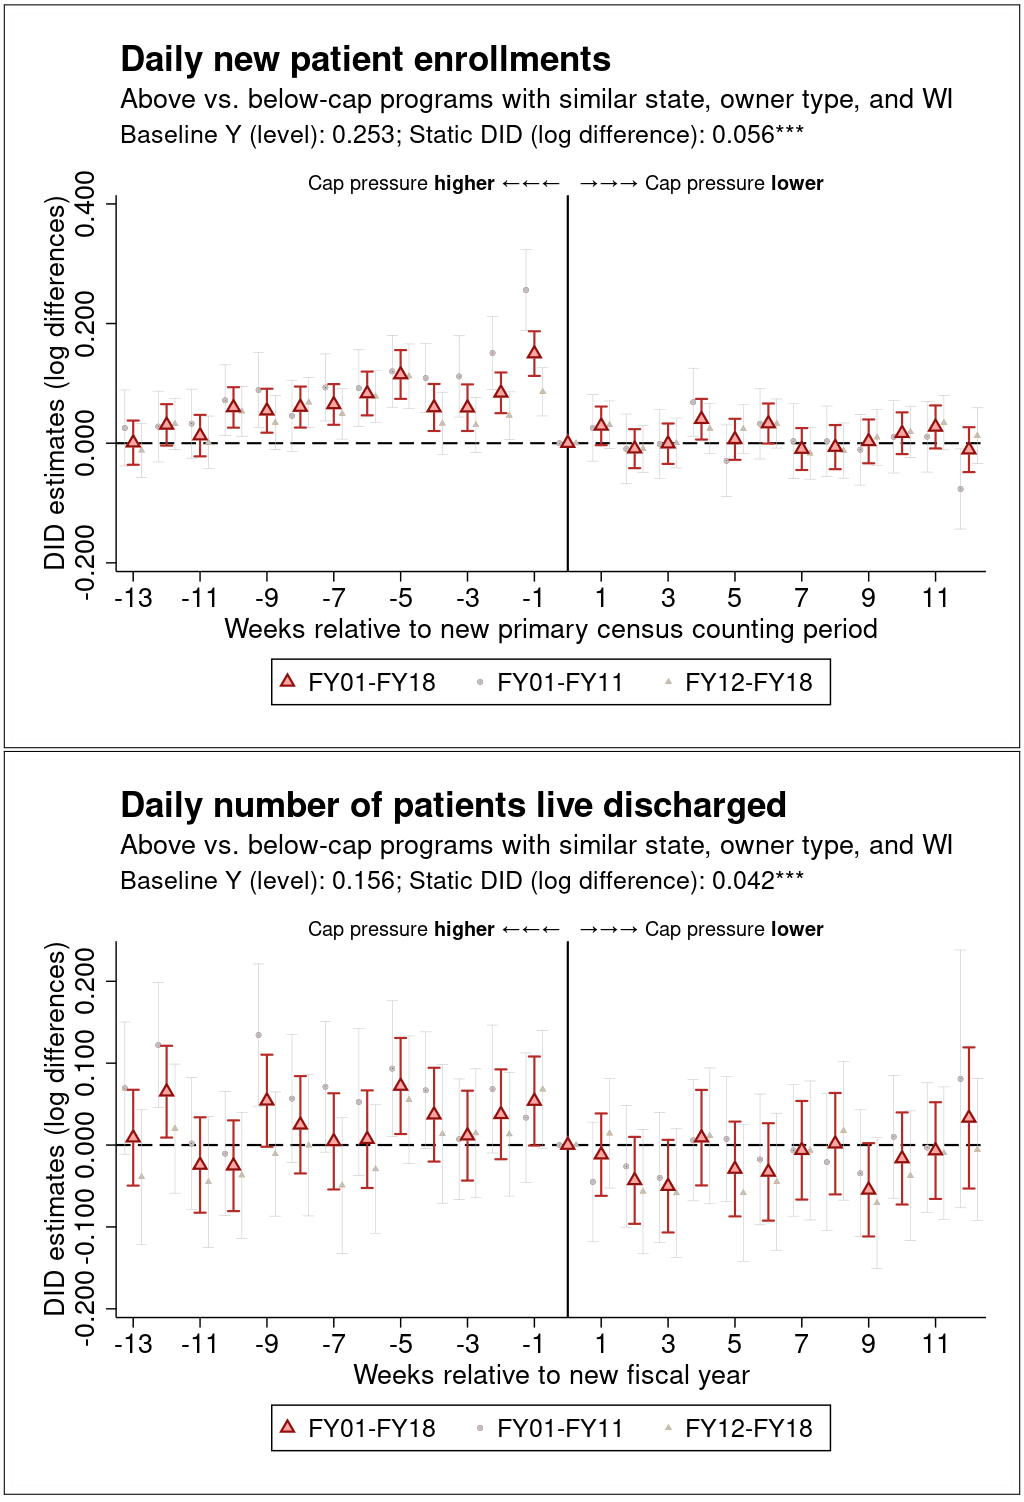


Fig. 4. This figure plots estimates of $(\beta_w : w \in W)$ from equation (6) in our DID analysis. The 95% confidence intervals are computed with program-level cluster-robust SEs. The headings report statistics estimated from the full sample period. Baseline Y is the average value of the outcome at relative time 0 among programs in $\mathcal{T}$. Static DID is an estimate of β from equation (5). The primary census counting periods are described in appendix A.

	Daily new enrollments (log diff.)		Daily live discharges (log diff.)	
	(1)	(2)	(3)	(4)
$D^{\text{Sta.}}_{ejw}$	0.056***	0.093***	0.042***	0.067***
	(0.007)	(0.012)	(0.010)	(0.016)
$D^{\text{Sta.}}_{ejw} \times 1[\text{Post2012}]$	-	-0.058***	-	-0.045**
	-	(0.014)	-	(0.020)
Estimator	PPML	PPML	PPML	PPML
Hospice FE	Y	Y	Y	Y
Period FE	Y	Y	Y	Y
B-line Y	0.253	0.253	0.156	0.156
N	1,407,530	1,407,530	1,372,130	1,372,130
Clusters	5,554	5,554	5,530	5,530

Tab. 2. This table reports our DID results. Program-level cluster robust SEs are in parenthesis. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Baseline Y is the average value of the outcome during relative time 0 among programs in $\mathcal{T}$. Dynamic estimates are plotted in figure 4.

Table 3 shows that programs on track to exceed the cap in an outgoing fiscal year differentially enroll patients who appear healthier on average. In particular, they have an 0.9 percentage point lower likelihood of having had a recent hospital stay, an 0.6 percentage point lower likelihood of having had a recent NH stay, and an 0.044 log point (4.5%) higher RDOL on average. They also appear less connected to hospice care: on average, they have an 0.020 log point (2.0%) longer average LLOS (despite having a 4.5% higher RDOL), 0.019 log point (1.9%) more lifetime live discharges, and 0.5 percentage point lower likelihoods of dying in hospice care.

Table 4 shows that programs on track to exceed the cap in an outgoing fiscal year live discharge patients who are largely similar to the inframarginal patients. We find that the they would take 0.048 log point (4.9%) more days to subsequently resume hospice care on average. But otherwise, we do not observe a statistically significant difference in their LOS in hospice care before the live discharge or their likelihood of subsequently resuming hospice care.

6.2.3 Characteristics of hospice care

We next estimated equations (5)-(6) with outcomes measuring other hospice care characteristics. Observations are weighted by patient volume where noted. We report the static DID estimates in table 5. We report the dynamic DID estimates in appendix figure A3.

First, table 5 shows that programs on track to exceed the cap in an outgoing fiscal year provided much the same rates of RHC, CHC, IRC, and GIC, suggesting that they do not adjust their levels

	LLOS (log Δ) (1)	Lifetime num. of programs (log Δ) (2)	Before enrolling, fraction with: hosp. stay (%point Δ) (3)	NH stay (%point Δ) (4)
$D^{\text{Sta.}}_{ejw}$	0.020***	0.003*	-0.009***	-0.006***
	(0.007)	(0.002)	(0.002)	(0.002)
Estimator	PPML	PPML	OLS	OLS
Hospice FE	Y	Y	Y	Y
Period FE	Y	Y	Y	Y
B-line Y	182.306	1.219	0.279	0.222
N	1,142,852	1,142,852	1,150,420	1,150,420
Clusters	5,534	5,534	5,540	5,540

	Fraction discharged in 30 days (%point Δ) (5)	RDOL (log Δ) (6)	Lifetime num. of live discharges (log Δ) (7)	Fraction who die in hospice (%point Δ) (8)
$D^{\text{Sta.}}_{ejw}$	0.001	0.044***	0.019**	-0.005***
	(0.001)	(0.009)	(0.008)	(0.002)
Estimator	OLS	PPML	PPML	OLS
Hospice FE	Y	Y	Y	Y
Period FE	Y	Y	Y	Y
B-line Y	0.089	396	0.604	0.866
N	1,150,420	1,141,473	1,123,808	1,138,993
Clusters	5,540	5,534	5,513	5,508

Tab. 3. This table reports our DID results for new enrollees. Program-level cluster robust SEs are in parenthesis. Observations are weighted by the number of new enrollees. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Baseline Y is the average of the outcome during relative time 0 among programs in $\mathcal{T}$. Dynamic estimates are plotted in figure A1.

	Prior LOS (log Δ) (1)	Days until resumption (if any) (log Δ) (2)	Fraction who later resume hospice: Anywhere (%point Δ) (3)	Same program (%point Δ) (4)	Elsewhere (%point Δ) (5)
$D^{\text{Sta.}}_{ejw}$	0.013	0.048**	-0.003	0.001	-0.004
	(0.008)	(0.019)	(0.003)	(0.003)	(0.003)
Estimator	PPML	PPML	OLS	OLS	OLS
Hospice FE	Y	Y	Y	Y	Y
Period FE	Y	Y	Y	Y	Y
B-line Y	285	245	0.703	0.310	0.393
N	676,718	498,173	682,363	682,363	682,363
Clusters	5,458	5,290	5,469	5,469	5,469

Tab. 4. This table reports our DID results for live discharged patients. Program-level cluster robust SEs are in parenthesis. Observations are weighted by the number of live discharged patients. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Baseline Y is the average of the outcome during relative time 0 among programs in $\mathcal{T}$. Dynamic estimates are plotted in figure A2.

of hospice care.[23] Second, it shows that they provided much the same rates of nurse visits, social worker visits, and HHA visits, suggesting that they do not alter the distribution of staff visits per patient to accommodate the additional enrollees.

	Fraction of patient-days with a visit by:				
	Nurse (%point Δ) (1)	HHA (%point Δ) (2)	Social worker (%point Δ) (3)		
β	0.000	0.001	-0.000		
	(0.001)	(0.001)	(0.000)		
Estimator	OLS	OLS	OLS		
Hospice FE	Y	Y	Y		
Period FE	Y	Y	Y		
B-line Y	0.219	0.324	0.040		
N	999,622	999,622	999,622		
Clusters	5,246	5,246	5,246		
	Fraction of patient-days with:				
	RHC (%point Δ) (4)	CHC (%point Δ) (5)	IRC (%point Δ) (6)	GIC (%point Δ) (7)	
β	-0.000	-0.000	-0.000	0.000*	
	(0.000)	(0.000)	(0.000)	(0.000)	
Estimator	OLS	OLS	OLS	OLS	
Hospice FE	Y	Y	Y	Y	
Period FE	Y	Y	Y	Y	
B-line Y	0.989	0.004	0.003	0.008	
N	1,419,834	456,274	974,974	987,220	
Clusters	5,558	3,384	4,214	4,045	

Tab. 5. This table reports our DID results for active patients. Program-level cluster robust SEs are in parenthesis. Observations are weighted by the number of active patients. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Baseline Y is the average of the outcome during relative time 0 among programs in $\mathcal{T}$. Dynamic estimates are plotted in figure A3.

6.3 Sensitivity and heterogeneity analyses

We conducted several sensitivity analyses. In particular, we examined whether our main results are sensitive to (1) a longer event window; (2) excluding state, ownership, and WI controls; and (3) excluding programs whose running average annual revenue was 90-100% of the cap at the start of the event windows. Appendix D.1 presents our results. The estimates are qualitatively similar. We also explored whether there is evidence of meaningful heterogeneity across profit-status, cap

[23]For each outcome, we excluded programs that provided no patient-days of that level of care in either the outgoing or incoming fiscal years. Our findings are not sensitive to their inclusion.

census, proximity to other hospice programs on track to exceed the cap, and affiliation with nursing homes. We report the results of these heterogeneity analyses in appendix D.2. There is suggestive evidence that hospice programs affiliated with nursing homes have more marginal live discharges at the end of an outgoing fiscal year, especially during the pre-2011 era. We did not find other evidence of significant heterogeneity along these dimensions.

6.4 Discussion

In sum, our DID analysis shows that hospice programs churn patients because of a cap on their average annual revenue. We find that programs on track to exceed the cap in an outgoing fiscal year increase enrollment and live discharge rates and we find evidence that a 2012 reform moderated this response. We also find evidence that the marginal enrollees are healthier on average and less connected to hospice care, which supports the idea that they have lower intrinsic demand for hospice care. Our heterogeneity analysis suggests while program characteristics such as ownership type may predict cap liabilities, they do not predict differential responses to the cap conditional on being on track to exceed the cap in an outgoing fiscal year.

However, our estimates also support the idea that hospice programs' adjustments to the cap are small relative to their cap liabilities. In particular, they imply that programs on track to exceed the cap in an outgoing fiscal year each have 1.54 additional enrollments and 0.37 additional live discharges throughout the outgoing fiscal year's last quarter on average. During 2001-2018, 1.54 additional enrollees would have raised a hospice program's MRR by at most $25,643 (in 2001) to $44,181 (in 2018), and 0.37 additional live discharges would have lowered its total payments by at most $3,391 (in 2001) to $5,960 (in 2018).[24] By contrast, programs that ultimately exceed the cap had a cap liability of $490,000 on average during 2001-2018.

These findings suggest that the cost of churning patients is high. We hypothesize that programs are primarily limited by non-pecuniary features of the Medicare hospice benefit. For instance, the

[24]With respect to enrollments, the cap was $16,651 (in 2001) and $28,689 (in 2018). If 1.54 additional patients were enrolled and had a LLOS of one day, then they would have raised their hospice program's MRR by $1.54 \times \$16,651 = \$25,643$ (in 2001) and $1.54 \times \$28,689 = \$44,181$ (in 2018). With respect to live discharges, daily payment rates for RHC at a program with a WI of 1 were $102 (in 2001), and $193 (in 2018, up to the first 60 days) or $151 (in 2018, after the 60th day). If 0.37 additional patients were live discharged 90 days before the end of the fiscal year and would have remained enrolled for 90 days but for their live discharge, then their live discharge would have reduced their program's annual revenue by $0.37 \times 90 \times \$102 = \$3,391$ (in 2001) and $0.37 \times (60 \times \$193 + 30 \times \$151) = \$5,960$ (in 2018).

Medicare hospice benefit requires that a non-hospice provider be consulted in enrollment and live discharge decisions, whenever possible. It also requires that Medicare beneficiaries who enroll in hospice care forgo coverage for their terminal illnesses. Each of these requirements raise the cost of enrolling and live discharging patients. Programs may also be limited by legal concerns. Gruber et al. (2023) study programs that have been sued under the federal False Claims Act (FCA) for enrolling patients with conditions that were not likely to lead to death within six months. While their analysis of these cases is not related to the cap, it suggests that a fear of litigation may deter programs facing cap liabilities from enrolling patients who are far from the eligibility margin.

7 Conclusion

We study how an average revenue cap can affect provider behavior and health care utilization and spending under capitation. In particular, we study a cap in the Medicare hospice benefit on hospice programs' average annual revenue. We find that programs on track to exceed the cap in an outgoing fiscal year increase enrollment and live discharge rates by 5.8% and 4.3% respectively during the fiscal year's fourth quarter. Characteristics of the marginal enrollees support the idea that the cap caused programs to move down the demand curve and enroll patients with a lower intrinsic demand for hospice care. We also find that programs that exceed the cap have an approximately $500,000 cap liability per year on average and that the wage index—a significant, visible, and administratively set determinant of payment rates—is nevertheless a significant predictor of cap liabilities.

In sum, our results suggests that a cap on average annual revenue can reduce health care spending in the market segment under the cap. Although hospice programs somewhat undercut the cap by churning their patients, they fell far short of minimizing their cap liabilities, which totaled approximately $185 million per year between 2001 and 2018. Their scope for inducing demand may be small because of non-pecuniary features of the Medicare hospice benefit. We believe that identifying the causal effect of the cap on market structure may be a productive area for future research.

8 Bibliography

1. Alexander, Diane. 2020. "How do doctors respond to incentives? Unintended consequences of paying doctors to reduce costs." Federal Reserve Bank of Chicago WP 2017-09.
2. Alexander, Diane and Molly Schnell. 2024. "The impact of physician payments on patient access, use, and health." *American Economic Journal: Economic Policy* 16(3): 142-77.
3. Anhang Price, Rebecca, Brian Stucky, Layla Parast, Marc N. Elliott, Ann Haas, Melissa Bradley, and Joan M. Teno. 2018. "Development of valid and reliable measures of patient and family experiences of hospice care for public reporting." *Journal of Palliative Medicine* 21(7): 924-32.
4. Arrow, Kenneth J. 1963. "Uncertainty and the welfare economics of medical care." *American Economic Review* 53 (5): 941-73.
5. Ata, Baris, Bradley L. Killaly, Tava Lennon Olsen, and Rodney P. Parker. 2012. "On hospice operations under Medicare reimbursement policies." *Management Science* 59 (5): 1027-44.
6. Braun, Robert Tyler, David G. Stevenson, and Mark Aaron Unruh. 2021. "Acquisitions of hospice agencies by private equity firms and publicly traded corporations." *JAMA Internal Medicine*, 181 (8): 1113-4.
7. Brot-Goldberg, Zarek C., Amitabh Chandra, Benjamin R. Handel, and Jonathan T. Kolstad. 2017. "What does a deductible do? The impact of cost-sharing on health care prices, quantities, and spending dynamics." *The Quarterly Journal of Economics* 132 (3): 1261-318.
8. Cattaneo, Matias D., Michael Jansson, Xinwei Ma. 2018. "Manipulation testing based on density discontinuity." *The Stata Journal* 18 (1): 234-61.
9. 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.
10. Chandra, Amitabh, David Cutler, and Zirui Song. 2012. "Who ordered that? The economics of treatment choices in medical care." In *Handbook of Health Economics Volume 2*, ed. Mark V. Pauly, Thomas G. McGuire, and Pedro P. Barros, 397-432. North Holland: Elsevier.
11. Clemens, Jeffrey, and Joshua D. Gottlieb. "Do physicians' financial incentives affect medical treatment and patient health?" *American Economic Review* 104 (4): 1320-49.
12. CMS. 2024. "Current Measures." Available online at https://www.cms.gov/medicare/quality/hospice/current-measures. Accessed in December 2024.
13. Correia, Sergio, Paulo Guimaraes, and Tom Zylkin. 2020. "Fast Poisson estimation with high dimensional fixed effects." *Stata Journal* 20 (1): 95-115.
14. Correia, Sergio. 2017. "Linear models with high-dimensional fixed effects: an efficient and feasible estimator." http://scorreia.com/research/hdfe.pdf. Accessed on October 2, 2023.
15. Dafny, Leemore. 2005. "How do hospitals respond to price changes?" *American Economic Review* 95(5): 1525-47.
16. Dalton, Christina Marsh, and W. David Bradford. 2019. "Better together: coexistence of for-profit and nonprofit firms with an application to the U.S. hospice industry." *Journal of Health Economics* 63: 1-18.

17. Decarolis, Francesco. 2015. "Medicare Part D: Are insurers gaming the low income subsidy design?" *American Economic Review* 105(4): 1547-80.
18. Dolin, Rachel, Laura C. Hanson, Sarah F. Rosenblum, Sally C. Stearns, George M. Holmes, Pam Silberman. 2017. "Factors driving live discharge from hospice: provider perspectives." *Journal of Pain and Symptom Management* 53 (6): 1050-6.
19. Dolin, Rachel, Pam Silberman, Denise A. Kirk, Sally C. Stearns, Laura C. Hanson, Donald H. Taylor Jr., and G. Mark Holmes. 2018. "Do live discharge rates increase as hospices approach their Medicare aggregate payment caps?" *Journal of Pain and Symptom Management* 55 (3): 775-84.
20. Einav, Liran, Amy Finkelstein, Sendhil Mullainathan, and Ziad Obermeyer. 2018. "Predictive modeling of US healthcare spending in late life." *Science* 360 (6396): 1462-65.
21. Einav, Liran, Amy Finkelstein, and Neale Mahoney. 2018. "Provider incentives and healthcare costs: evidence from long-term care hospitals." *Econometrica* 86 (6): 2161-219.
22. Einav, Liran, Amy Finkelstein, Yunan Ji, and Neale Mahoney. 2022. "Voluntary regulation: evidence from Medicare payment reform." *Quarterly Journal of Economics*: 565-618.
23. Eliason, Paul J., Benjamin Heebsh, Riley J. League, Ryan C. McDevitt, and James W. Roberts. 2022. "The effect of bundled payments on provider behavior and patient outcomes." Accessed on November 5, 2022.
24. Eliason, Paul J., Paul L. E. Grieco, Ryan C. McDevitt, and James W. Roberts. 2018. "Strategic patient discharge: the case of long-term care hospitals." *American Economic Review* 108 (11): 3232-65.
25. Ellis, Randall P., and Thomas G. McGuire. 1986. "Provider behavior under prospective reimbursement: cost sharing and supply." *Journal of Health Economics* 5: 129-151.
26. Friedberg, Mark W., Peggy G. Chen, Chapin White, Olivia Jung, Laura Raaen, Samuel Hirshman, Emily Hoch, Clare Stevens, Paul B. Ginsburg, Lawrence P. Casalino, Michael Tutty, Carol Vargo, and Lisa Lipinski. 2015. "Effects of health care payment models on physician practice in the United States." Santa Monica, CA: RAND Corporation. https://www.rand.org/pubs/research_reports/RR869.html.
27. Furuno, Jon P., Brie N. Noble, Mary L. McPherson, Kate L. Lapane, Leah Sera, Shigeko Izumi, and Jennifer Tjia. 2020. "Variation in hospice patient and admission characteristics by referral location." *Medical Care* 58(12): 1069-74.
28. Gandhi, Sabina Ohri. 2012. "Differences between non-profit and for-profit hospices: patient selection and quality." *International Journal of Health Care Finance and Economics* 12(2): 107-27.
29. Gardner, John. 2021. "Two-stage difference in differences." https://jrgcmu.github.io/2sdd_current.pdf. Accessed on October 15, 2021.
30. Goodman-Bacon, Andrew. 2021. "Difference-in-differences with variation in treatment timing." *Journal of Econometrics* 225(2):254-77.
31. Grabowski, David C., Christopher C. Afendulis, and Thomas G. McGuire. 2011. "Medicare prospective payment and the volume and intensity of skilled nursing facility services." *Journal of Health Economics* 30: 675-84.

32. Gruber, Jonathan, David H. Howard, Jetson Leder-Luis, and Theodore L. Caputi. 2023. "Dying or lying? For-profit hospices and end of life care." NBER Working Paper 31035.

33. Gupta, Atul. 2021. "Impacts of performance pay for hospitals: the readmissions reduction program." *American Economic Review* 111 (4): 1241-83.

34. Gupta, Atul, Ambar La Forgia, and Adam Sacarny. 2024. "Turbocharging profits? Contract gaming and revenue allocation in healthcare." NBER Working Paper 32564.

35. Harrison, Krista L., Irena Cenzer, Claire K. Ankuda, Lauren J. Hunt, and Melissa D. Aldridge. 2022. "Hospice improves care quality for older adults with dementia in their last month of life." *Health Affairs* 41(6): 821-30.

36. Ho, Kate and Ariel Pakes. 2014. "Hospital choices, hospital prices, and financial incentives to physicians." *American Economic Review* 104(12): 3841-84.

37. Jenkins, Todd M., Kathryn L. Chapman, Christine S. Ritchie, Donna K. Arnett, Gerald McGwain, Stacy Cofield, and H. Michael Maetz. 2011. "Barriers to hospice care in Alabama: provider-based perceptions." *American Journal of Hospice & Palliative Medicine* 28 (3): 153-60.

38. Kelley, Amy S., Partha Deb, Qingling Du, Melissa D. Aldridge Carlson, and R. Sean Morrison. 2013. "Hospice enrollment saves money for Medicare and improves care quality across a number of different lengths-of-stay." *Health Affairs* 32 (3): 552-61.

39. Kofman, Ava. 2022. "Endgame: How the visionary hospice movement became a for-profit hustle." *ProPublica* and *The New Yorker*: November 28. https://www.newyorker.com/magazine/2022/12/05/how-hospice-became-a-for-profit-hustle.

40. Liebman, Jeffrey B., and Neale Mahoney. 2017. "Do expiring budgets lead to wasteful year-end spending? Evidence from federal procurement." *American Economic Review* 107(11): 3510-49.

41. Lee, David W. and Frank Levy. 2012. "The sharp slowdown in growth of medical imaging: an early analysis suggests combination of policies was the cause." *Health Affairs* 31 (8): 1876-84.

42. McCrary, Justin. 2008. "Manipulation of the running variable in the regression discontinuity design: a density test." *Journal of Econometrics* 142 (2): 698-714.

43. Medicare Payment Advisory Commission (MedPAC). 2010. *Report to Congress*. Medicare Payment Advisory Commission (MedPAC). Washington, DC, March.

44. MedPAC. 2020. *Report to Congress*. Medicare Payment Advisory Commission (MedPAC). Washington, DC, March.

45. MedPAC. 2022. *Report to Congress*. Medicare Payment Advisory Commission (MedPAC). Washington, DC, March.

46. Meltzer, David, Jeanette Chung, and Anirban Basu. 2002. "Does competition under Medicare prospective payment selectively reduce expenditures on high-cost patients?" *RAND Journal of Economics* 33(3): 447-68.

47. Norton, Edward C., Jun Li, Anup Das, and Lena M. Chen. 2018. "Moneyball in Medicare." *Journal of Health Economics* 61: 259-73.

48. Office of Inspector General (OIG). 2021. "Palmetto GBA, LLC, accurately calculated hospice cap amounts but did not collect all cap overpayments." Report No. A-06-19-08003.

49. Office of Inspector General (OIG). 2022a. "National Government Services, Inc., accurately calculated hospice cap amounts but did not collect all cap overpayments." Report No. A-06-21-08004.

50. Office of Inspector General (OIG). 2022b. "Medicare hospice provider compliance audit: Vitas Healthcare Corporation of Florida." Report No. A-02-19-01018.

51. Park, Sungchul, Anirban Basu, Norma Coe, and Fahad Khalil. 2017. "Service-level selection: strategic risk selection in Medicare Advantage in response to risk adjustment." NBER Working Paper 24038.

52. Pauly, Mark V. 2000. "Insurance reimbursement." *Handbook of Health Economics* 1: 537-60.

53. Plotzke, Michael, T.J. Christian, Elizabeth Axelrod, Meaghan Hunt, and Allison Muma. 2015. "Medicare hospice payment reform: analysis of how the Medicare hospice benefit is used." Abt. Associates Technical Report HHSM-502005-00018I.

54. Plotzke, Michael, Thomas Christian, Kim Groover, Zinnia Harrison, Ihsan Abdur-Rahman, and Cindy Massuda. 2021. "Construction and performance of the hospice care index claims-based quality measure." *Innovations in Aging* 5(S1).

55. Polsky, Daniel, Guy David, Jianing Yang, Bruce Kinosian, and Rachel Werner. 2014. "The effect of entry regulation in the health care sector: the case of home health." *Journal of Public Economics* 110: 1-14.

56. Rao, Jordan. 2011. "Concerns about costs rise with hospices' use." *The New York Times*, June 27. https://www.nytimes.com/2011/06/28/health/28hospice.html.

57. Rosenkranz, David A. 2024. "Entry barriers in provider markets: evidence from dialysis certificate-of-need programs." https://davidrosenkranz.github.io/davidrosenkranz.com/esrd-con.pdf. Accessed on December 1, 2024.

58. Rosenkranz, David A., Lindsay White, Chuxuan Sun, Katherine E.M. Miller, and Norma B. Coe. 2024. "Market segmentation by profit status: evidence from hospice." *Health Affairs Scholar*. https://doi.org/10.1093/haschl/qxae160.

59. Sack, Kevin. 2007. "In hospice care, longer lives means money lost." *The New York Times*: November 27. https://www.nytimes.com/2007/11/27/us/27hospice.html.

60. Stevenson, David G., Jesse B. Dalton, David C. Grabowski, and Haiden A. Huskamp. 2015. "Nearly half of all Medicare hospice enrollees received care from agencies owned by regional or national chains." *Health Affairs; Chevy Chase* 34 (1): 30-8.

61. Stevenson, David, Nicholas Sinclair, Shiyuan Zhang, Laurie Meneades, and Haiden Huskamp. 2018. "Trends in nursing home-hospice contracting and common ownership between hospice agencies and nursing homes: final report." U.S. Department of Health and Human Services, Washington, DC. November 16.

62. Teno, Joan M., Pedro L. Gozalo, Ian C. Lee, Sylvia Kuo, Carol Spence, Stephen R. Connor, and David J. Casarett. 2011. "Does hospice improve quality of care for persons dying from dementia?" *Journal of the American Geriatric Society* 59 (8): 1531-1536.

63. Teno, Joan, Michael Plotzke, Pedro Gozalo, and Vincent Mor. 2014. "A national study of live discharges from hospice." *Journal of Palliative Medicine* 17 (10): 1121-7.

64. Waldman, Peter. 2012. "Hospices dumping patients while escaping millions owed to U.S." *Bloomberg Business*. February 3. https://www.bloomberg.com/news/articles/2012-02-03/hospices-dump-patients-while-escaping-multimillion-dollar-taxpayer-refunds.

65. Wooldridge, Jeffrey M. 1999. "Distribution-free estimation of some nonlinear panel data models." *Journal of Econometrics* 90 (1): 77-97.

66. Zuckerman, Rachel B., Sally C. Stearns, and Steven H. Sheingold. 2015. "Hospice use, hospitalization, and Medicare spending at the end of life." *The Journals of Gerontology: Series B* 71 (3): 569-80.

APPENDIX

A Application of the payment cap

A.1 What determines the cap?

The cap is set by CMS each year. Prior to and including 2015, the cap varied annually with the medical expenditure category of the Consumer Price Index for urban consumers published by the Bureau of Labor Statistics. Since 2016, the cap has varied annually with hospice payment rates.[25] The cap has grown from $6,500 at the inception of the hospice benefit to $32,487 in FY 2023.[26]

Within a given fiscal year, the cap varies across hospice programs to a limited extent. The first fiscal year assigned to a new hospice program is the first fiscal year ending at least twelve months after the hospice program received its Medicare certification. New hospice programs' first fiscal years may therefore span parts of two fiscal years. For these hospice programs, the cap is a weighted average of the two fiscal years' caps. The weights are the fraction of days that the hospice program was operating in each of the two fiscal years. The cap does not vary within a fiscal year along any other dimensions.[27]

A.2 What is a fiscal year?

For every fiscal year t through FY 2016, the fiscal year began on November 1 of calendar year $t-1$ and ended on October 31 of calendar year t. In FY 2017, the fiscal year began on November 1, 2016 and ended on September 30, 2017. Since FY 2018, each fiscal year t has begun on October 1 of calendar year $t-1$ and ended on September 30 of calendar year t.

A.3 How are hospice programs' average annual payments per patient computed?

In each fiscal year, each hospice program's average annual payment per patient is the program's gross Medicare payment divided by its cap census. The cap census accounts for the possibility that each Medicare beneficiary may utilize the hospice benefit across fiscal years and hospice programs.

[25] Federal Register, vol. 80(151): 47207.

[26] The cap was initially 40% of Medicare's average expenditures for cancer patients in their last six months of life, or $6,500. Since 1983, the cap has been annually increased to keep pace with inflation in medical expenditures and baseline hospice payment rates. See Federal Register, vol. 48(163):38156-7 and Federal Register vol. 80(151):47147.

[27] Medicare hospice transmittal 156.

A.4 Payments

Each hospice program's total payment in each fiscal year is the sum of the per diem payments made to that program for patient-days that fall within that fiscal year, irrespective of the day the payments were actually remitted. Mathematically, let j index hospice programs, let t index fiscal years, let i index beneficiaries, let d index days, let $a(i, j, d) := 1[i \text{ was enrolled at } j \text{ on } d]$, let $t(d)$ be the fiscal year corresponding to d, and let $\text{PD}(i, j, d)$ be the per diem payment made to j for beneficiary i's care on d. Then for every (j, t), total payments were:

$$\text{Payment}_{jt} := \sum_{d:t(d)=t} \sum_{i:a(i,j,d)=1} \text{PD}(i, j, d) \tag{A.1}$$

The per diem rate is administratively set by CMS according to a formula that varies with beneficiary, day, and program. For a given (i, j, d), the per diem rate consists of a base payment, a so-called "labor share," and a wage index adjustment.

The base payment for a given patient-day depends on the level of hospice care the patient received on that day, the patient's previous hospice utilization, the patient's program's compliance with Medicare's quality reporting requirements, and the day's fiscal year. There are four treatment types: RHC, CHC, IRC, and GIC. Their base payments are annually adjusted by Medicare to account for growth in the hospital market basket. Within each fiscal year, the base payment is generally lowest for RHC, then for IRC, then for GIC, then for CHC. Since January 1, 2016, the RHC base payments have also been higher for each patient's first 60 days in hospice care than they have been for their subsequent days. Finally, since FY 2014, CMS has reduced base payment rates by 2% for hospice programs that have not complied with certain quality reporting requirements.[28]

The base payment for each patient-day is divided into a "labor amount" and a "non-labor amount" using an administratively set "labor share." The labor share varies according to treatment type. It is 0.6871 for RHC, 0.6871 for CHC, 0.5413 for IRC, and 0.6401 for GIC, and has not changed since at least FY 2001. The labor amount is adjusted by the wage index to compensate hospice programs operating in geographic markets with higher labor costs. Each program's wage index value varies from year-to-year with how labor costs at its CBSA compare to nationwide labor costs.

The base payments, labor shares, and wage index adjustments are combined to calculate the per diem payment rates as follows. Let $k(i, d) \in \{\text{RHC}, \text{CHC}, \text{IRC}, \text{GIC}\}$ indicate whether beneficiary i received RHC, CHC, IRC, or GIC on day d. Let $\text{BP}_k(i, j, d)$ be the beneficiary's base payment rate on day d at program j for treatment type k. (Recall that for a given (j, d), $\text{BP}_k(i, j, d)$ does not vary with i unless $k = \text{RHC}$.) Let LS_k be the labor share for treatment type k. Finally, let $\text{WI}(j, t)$

[28] Federal Register, vol. 80(151): 47207.

be the wage index adjustment for program j in fiscal year t. Then the per diem rates were:

$$\text{PD}(i,j,d) := \sum_k 1[k(i,d)=k]\Big(\text{BP}_k(i,j,d)(1-LS_k) + \text{BP}_k(i,j,d)\text{LS}_k\text{WI}(j,t(d))\Big) \quad \text{(A.2)}$$

A.5 Cap census

Since FY 1984, patients have been counted toward cap censuses with one of two formulas, or "counting methods:" the streamlined method and the proportional method. The particular method used for each patient depends on that patient's program and the fiscal year.

Under the streamlined method, patients who were only ever affiliated with one hospice program were counted toward that program's cap census on their first day enrolled in hospice. A patient was counted toward their program's cap census in fiscal year $t \in \{2001, ..., 2016\}$ if that day fell between September 28 of calendar year $t-1$ and September 27 of calendar year t. They were counted toward their program's cap census in FY 2017 if that day fell between September 28, 2016 and September 30, 2017. Finally, since FY 2018, they have been counted toward their program's cap census in a fiscal year t if that day fell between October 1 of calendar year $t-1$ and September 30 of calendar year t. Patients who were ever affiliated with more than one program were counted toward each program's cap census proportionally. Specifically, for each fiscal year t that such a patient was enrolled in hospice and for each program j where the patient was ever enrolled, that patient's contribution to (j,t)'s cap census is that fraction of their total number of patient-days that they spent at (j,t).[29] Under the proportional method, this formula was applied for all patients—not just those who were ever affiliated with multiple programs.

All hospice programs used the streamlined method until FY 2011. (Hence, we sometimes refer to this as the "pre-2011 counting method.") Around that time, hospice programs appealed to contest the streamlined method. In response to the appeal, CMS announced in 2011 that hospice programs could elect by October 1, 2011 to have their cap censuses for fiscal years prior to 2011 recalculated using the proportional method. If a hospice program made this election, then its cap censuses for fiscal years 2012 and onward were also made using the proportional method.[30] (Hence, we sometimes refer to this as the "post-2012 counting method.")

All other programs would also be switched to the proportional method for fiscal years 2012 and onward unless they made an election to continue using the streamlined method shortly after the end of FY 2012. If they did not make this election and therefore began using the proportional method in FY 2012, then their pre-2011 cap censuses were unchanged, even if they served beneficiaries in FY

[29] If a beneficiary is enrolled with two hospice programs on the same say, then that day counts toward both program's cap censuses and raises the beneficiary's total number of beneficiary-days by two. For example, this may occur on days when beneficiaries are switching from one hospice program to another. See Medicare hospice transmittal 156.

[30] Medicare hospice transmittal 156.

2012 or later who had already been counted in a pre-2011 fiscal year using the streamlined method. Consequently, some beneficiaries' lifetime contributions to hospice programs' cap censuses are more than one.[31]

Those programs that elected to continue using the streamlined method after FY 2011 could elect to switch to the proportional method at any time. Once they elected to switch, their cap censuses would be calculated using the proportional method in the year that they elected to switch and in all subsequent fiscal years. CMS would also recalculate their cap censuses using the proportional method for up to three fiscal years prior to the fiscal year that they elected to switch. Once a hospice program begins using the proportional counting method, it cannot switch back. Hospice programs that received their Medicare certification on or after October 1, 2011 use the proportional method.[32]

Only 486 hospice programs used the streamlined method in FY 2013.[33] This has weakly decreased since 2013, but data identifying which programs used which methods in which fiscal years are not available. Consequently, we estimate cap liabilities assuming that all hospice programs used the streamlined method during FY 2001 - FY 2011 and the proportional method thereafter. Since most patients are only ever affiliated with one hospice program, we define the primary census counting period for each fiscal year $t \in \{2001, ..., 2011\}$ to be September 28 of calendar year $t-1$ to September 27 of calendar year t; and we define it for each fiscal year $t \in \{2012, ..., 2018\}$ to line up with the fiscal year's start and end dates.

In any given fiscal year t prior to and including FY 2016 when a hospice program switched from the streamlined method in FY $t-1$ to the proportional method in FY t, patient-days between September 28 of calendar year $t-1$ and the October 31 of calendar year $t-1$ were not counted toward any cap census. For example, consider a patient who first enrolled in hospice on September 28, 2011 at a program that was switching from the streamlined method to the proportional method between FY 2011 and FY 2012. That patient would not have been counted toward the FY 2011 cap census because the window to be counted using the streamlined method had closed on September 27, 2011. If they passed away on October 31, 2011, then they would not have been counted toward the program's FY 2012 cap census either because they passed away before FY 2012 began on November 1, 2011. If they instead passed away on November 30, 2011, then their patient-days between November 1 and 30 would have been counted toward the FY 2012 cap census in proportion to their total length-of-stay, which would have included the 34 days between September 28 and October 31. However, those 34 days would not have themselves contributed to either fiscal year's cap census. Consequently, some patients' lifetime contributions to hospice programs' cap censuses are less than one.[34]

[31] Medicare hospice transmittal 156.
[32] Medicare hospice transmittal 156.
[33] Federal Register, vol. 80(151): August 6, 2015
[34] Medicare hospice transmittal 156.

Under both the streamlined and proportional methods, patients' contributions to their hospice programs' cap censuses depend on their lifetime hospice utilization. In practice, for any given fiscal year t, CMS may retroactively adjust hospice programs' FY t cap censuses with updated utilization data collected up to fiscal year $t+3$. In each fiscal year t, hospice programs must repay an estimate of their cap libailities within five months of the end of t.[35] Their cap liabilities will be weakly larger than this estimate as some of their patients continue to use hospice care through year $t+3$. Hospice programs may be subsequently asked to repay the difference.[36] Therefore, our definitions of cap-related variables such as Liability_{jt}, Census_{jt}, and PPP_{jt} are with respect to this three-year look-back period.

A.6 Assigning (program, days) to fiscal years

We linked each (program, day) realization observed in the hospice claims data to its corresponding fiscal year. This link is usually based which fiscal year's start and end dates contain each day. For instance, days between November 1, 2007 and October 31, 2008 (inclusive) are linked to fiscal year 2008. However, a new hospice program's first fiscal year is that fiscal year ending at least twelve months after the program received its Medicare certification. For instance, for hospice programs that received their Medicare certifications between November 2, 2006 and November 1, 2007 (inclusive), days in this range are part of FY 2008.

We identified hospice programs' Medicare certification dates as follows. In the Medicare provider of service files ("POS files," hereafter), we identified each program's earliest observed Medicare certification date. For each hospice program j, let this date be $\text{FirstCertDay}(j)$. Similarly, let the earliest observed date associated with any of program j's Medicare claims be $\text{FirstClaimDay}(j)$. For each program j such that $\text{FirstCertDay}(j)$ is non-missing and earlier than $\text{FirstClaimDay}(j)$, we treated $\text{FirstCertDay}(j)$ as the program's Medicare certification date. But for some programs j, $\text{FirstCertDay}(j)$ is either missing or greater than $\text{FirstClaimDay}(j)$. (This represents approximately 45% of hospice programs.) We assume that this indicates an error in the POS files. For such programs j, if $\text{FirstClaimDay}(j)$ is strictly greater than January 1, 2000, then we treated $\text{FirstClaimDay}(j)$ as the program's Medicare certification date. (This represents approximately 20% of hospice programs.) Otherwise, we assume that the first fiscal year ending 12 months after program j's Medicare certification is 2000 or earlier. (This calculation involves the date January 1, 2000 because our sample of the Medicare claims data begins in calendar year 2000, so $\text{FirstClaimDay}(j)$ is January 1, 2000 for many programs.)

[35] See 42 CFR 418.308.

[36] Federal Register, Volume 80(151): August 6, 2015. See also the CMS Manual System: Pub 100-02 Medicare Benefit Policy, transmittal 156, dated June 1, 2022.

A.7 Sequestration adjustment

The Budget Control Act of 2011 triggered a sequestration reduction to Medicare spending which cut hospice payments made on or after April 1, 2013 by 2 percent. In response, CMS applied a sequestration adjustment to the cap liability formula. First, CMS computed programs' payments but for the 2% sequestration reduction. Second, CMS compared this pre-sequestration amount to the cap to compute a pre-sequestration cap liability. Third, CMS reduced the pre-sequestration cap liability by the percentage difference between hospice programs' actual payments and their payments but for the sequestration reduction.[37] After FY 2013—when payments were only reduced for part of the fiscal year—CMS's sequestration adjustment is equivalent to reducing the cap by 2%.

B Data

B.1 Discharge codes

Each hospice claim is associated with one of several discharge codes. We grouped the discharge codes into five categories: discharged home, discharged to another hospice, discharged to another health care institution, discharged to an unknown location (collectively, "discharged alive") and discharged dead. We validated the discharge codes in three steps:

1. For each claim, if the patient was observed in the same hospice on the day after the claim's last day, then we disregarded the discharge code on that claim. If they were observed in a different hospice on the day after the claim's last day, then we treated that claim as ending in a code for "discharged to another hospice." If the claim's last day coincided with the patient's date of death in the Medicare Denominator file, then we treated that claim as ending in a code for "discharged dead."

2. Otherwise, if a claim ended in a code for "discharged to another hospice," but the patient was not observed in any hospice on the next day, then we treated that claim as ending in a code for "discharged to an unknown location."

3. Otherwise, if a claim ended in a code for "discharged dead," but the patient's date of death in the Medicare Denominator file was after the last day of the claim, then we treated that claim as ending in a code for "discharged to an unknown location."

After we assigned a discharge code to each claim, we assigned to each (program, day, patient) observation all discharge codes from claims associated with that (program, patient) and ending on

[37] Opinion of Judge Daniel A. Bress in the matter of *Silverado Hospice, Inc. vs. Xavier Becerra*, filed August 1, 2022.

that day. There were cases where multiple discharge codes for "discharged to another institution," "discharged home," and "discharged to an unknown location" were each assigned to the same (program, day, patient) observation because the (program, patient) had multiple claims ending on that day. In such cases, we eliminated all but one discharge code using the following priority ranking: "discharged to another institution," then "discharged home," and then "discharged to unknown." (By construction, no (program, day, patient) observation has multiple discharge codes if any of the codes indicate "discharged dead" or "discharged to another hospice.")

B.2 Patient demographics

We relied on the patient demographic data in the 2000-2019 Medicare Denominator files. In particular, we identified each beneficiary's date of birth, race, sex, and date of death. It is not common for one beneficiary to have multiple distinct realizations of these variables. When that happens, we assigned the beneficiary to whichever realization was reported first. In calculating hospice patients' remaining lifetimes after enrolling in hospice, we assume that patients with missing dates of death survived until December 31, 2019. (We re-estimated our DID analysis for RDOL including only decedents. The results were qualitatively similar.)

We also relied on the patient residence data in the hospice claims. In particular, we identified for each (program, day, patient) observation the corresponding ZIP code reported in claims associated with that observation. If a single (program, day, patient) observation was ever associated with multiple ZIP codes, we treated the ZIP code as missing for that observation.

B.3 Hospice program business records

We relied on hospice programs' business data in the 2000-2019 Medicare provider-of-service files. In particular, we identified each program's ownership type (i.e., for-profit, not for-profit, public, or other/unknown), county, state, and facility name. We also used the programs' original participation dates and their earliest observed hospice claims to determine their ages.

We also used the hospice claims data to identify the first and last dates a hospice program's CCN is associated with a hospice claim. See appendix A.6 for details about how we used the claims data to identify when hospice programs began operating. We say that a hospice program's CCN is terminated in a given fiscal year if it is not associated with any hospice claims in the next fiscal year. CCNs are terminated when their underlying provider agreements are terminated (Medicare State operations Manual §2779). Provider agreements may be voluntarily terminated when a provider exits, when a provider goes bankrupt (Medicare Financial Management Manual, Chapter 3, §140.6.2), or when an institutional provider (such as a hospice program) undergoes a change of ownership and the new owner does not accept the institutional provider's existing provider agreement (Medicare

State Operations Manual §3210.5A). Provider agreements may also be involuntarily terminated for several reasons, including because of a failure to make a "satisfactory overpayment arrangement[]," (Medicare State Operations Manual §3028D). To the best of our knowledge, there have been no systematic taxonomies of CCN terminations in the U.S. hospice industry, so we cannot distinguish exits from other events.

B.4 Wage index, daily payment rates, and the cap

We relied on wage index, daily payment rates, and cap data published by Medicare in its annual transmittals to hospice programs and in the Federal Register. Between FY 2001 and FY 2018, the wage index data were reported alternately at the MSA-level, CBSA-level, or county-level. We linked the MSAs and CBSAs with counties to enable linking the wage index data with hospice programs.

B.5 Hospitalizations

We identified hospitalizations using the 2000-2019 MedPAR files. We linked the providers in the MedPAR files with their records in the provider-of-service files to determine which providers were hospitals and skilled nursing facilities. We determined days that hospice patients spent in a hospital and skilled nursing facility using the admission and discharge dates associated with each claim in the MedPAR files. We created a patient-level variable indicating whether each hospice patient was hospitalized within 7-days of their hospice enrollment date.

B.6 Nursing home stays

We constructed nursing home spells using the Minimum Dataset ("MDS," hereafter) and MedPAR. The MDS contains records of Medicare beneficiaries' health assessments generated when they are admitted to nursing homes ("entry assessment," hereafter), discharged from nursing homes ("discharge assessment," hereafter), and at regular intervals during their nursing home stays ("ongoing assessment," hereafter). We used these assessments to create a (patient, nursing home, day)-level dataset. Each observation corresponds to a day when an assessment was made for a patient by a nursing home. We excluded some difficult-to-parse patients and observations. For instance, we excluded assessments if that assessment's nursing home was neither the patient's previous assessment's nursing home or their subsequent assessment's nursing home. We also excluded patients if they were ever observed to have assessments from three nursing homes on the same day.

We constructed nursing home spells from the (patient, nursing home, day) data by connecting consecutive assessment dates. A patient's spell at a given nursing home began on (1) the day of

the patient's first-observed assessment; (2) when an entry assessment was observed at that nursing home; (3) when an ongoing assessment was observed at that nursing home and the previous observation was an assessment at another nursing home; or (4) when an ongoing assessment was observed at that nursing home and the previous observation was a discharge assessment at the same nursing home. A patient's spell at a given nursing home ended on (1) the day of the patient's last-observed assessment; (2) when a discharge assessment was observed at that nursing home; (3) when an ongoing assessment was observed at that nursing home and the next observation was an assessment at another nursing home; or (4) when an ongoing assessment was observed at that nursing home and the next observation was an entry assessment at the same nursing home. We joined (patient, day)-level nursing home stay records derived from the MDS with (patient, day)-level skilled nursing facility stay records in the MedPAR files. We created a patient-level variable indicating whether each hospice patient experienced a nursing home stay within 7-days of their hospice enrollment date.

B.7 Hospice staff visits

We identified days that a hospice staff visit occurred using the hospice claims data. In particular, we used the revenue center codes associated with each claim to identify skilled nurse, social worker, and home health aide visits associated with each claim-day. These visits were not reported from 2000 to 2007, so we exclude these years of data in analyses related to staff visits.

B.8 Measuring each program-year's proximity to other program's that exceeded the cap in that year

We constructed a measure of each program-year (j,t)'s proximity to other programs that exceeded the cap in that year using the extent of overlap of hospice programs' patient populations. For each ZIP code z in each fiscal year t, define the set of hospice programs serving patients in (z,t) as J_{zt}. For each program j, define I_{jzt} as the set of patients living in z and being treated in (j,t). For each (j,z,t) such that $j \in J_{zt}$, define the average fraction of other hospice programs that operate in (z,t) and exceed the cap in FY t weighted by their share of (z,t)'s hospice patient population as:

$$D_{jzt} := \sum_{j' \in J_{zt} \setminus \{j\}} 1[\text{Over the cap}]_{j't} \left(\frac{|I_{j'zt}|}{\sum_{j' \in J_{zt} \setminus \{j\}} |I_{j'zt}|} \right) \tag{B.3}$$

Then for each (j,t), define the average of D_{jzt} weighted by the share of (j,t)'s patient volume across (z,t) as:

$$D_{jt} := \sum_{z:j\in J_{zt}} D_{jzt}\left(\frac{|I_{jzt}|}{\sum_{z:j\in J_{zt}}|I_{jzt}|}\right) \tag{B.4}$$

This is a measure of proximity in the sense that it is increasing in the extent of geographic overlap between (j,t)'s patients and patients enrolled in hospices that exceeded the cap.

B.9 Levels of hospice care

We identified the level of hospice care associated with each patient-day using the claims data. Each claim includes several revenue center lines reporting that for a span of days, a number of units of RHC, CHC, GIC, or IRC were provided on those days. For RHC, GIC, and IRC, the units are measured in days. For CHC, the units are measured in hours (prior to and including 2006) or fifteen minute increments (since 2007). We converted the CHC units to days. For most claims, the dates and units on each revenue center line are sufficient to determine the level of hospice care provided on each claim-day. For instance, for claims spanning seventeen days and reporting 17 units of RHC, we say that one unit of RHC was provided on each day. Similarly, for claims spanning seventeen days and reporting 14 units of RHC on day 1 and 3 units of GIC on day 15, we say that one unit of RHC was provided on each of the first fourteen days and one unit of GIC was provided on the final three days. We distributed the units of each level of hospice care evenly across a given claim's days if we could not determine which levels of hospice care were provided on which day. We aggregated the (claim, day)-level data to the (provider, day, patient)-level by summing the number of units of RHC, CHC, GIC, or IRC provided to each patient by each program on each day across all claims.

C Additional descriptive statistics

The cap liability statistics we report in section 5 are consistent with those reported elsewhere. For instance, MedPAC (2020) reports that in 2002, average cap liabilities were $470,000 among 2.6% of programs that exceeded the cap. We estimate that they were $471,000 among 3.4% of programs that exceeded the cap. They also estimate that in 2014, cap liabilities were $370,000 among 12.1% of programs that exceeded the cap. We estimate that they were $390,000 among 13.9% of programs. These differences are attributable in part to differences in our calculation methodologies. MedPAC assumed that each patient's cap census contribution in each fiscal year is updated to account for that patient's hospice utilization in the following 14-15 months. We assumed that it is updated to

account for that patient's hospice utilization in the following three years to reflect that CMS may reopen a given program-year's cap liability during the following three fiscal years.

Our cap liability estimates indicate that cap risk is geographically clustered and autocorrelated. Figure A4 plots state-level trends in the average annual fraction of hospice programs that exceeded the cap. It shows that between 2001 and 2009, a larger fraction of hospice programs in MS, AL, AZ, and OK exceeded the cap than in other states. It also shows that since 2010, an increasing fraction of hospice programs in California have exceeded the cap. Figure A5 plots for the set of programs that exceeded the cap at least once the association between (1) the number of fiscal years between FY 2001 and FY 2018 that they were operating and (2) the average fraction of those years that they exceeded the cap. Programs that exceeded the cap at least once exceeded it in 49% of their observed years of operation between FY 2001 and FY 2018 on average.

We also investigate whether any individual hospice program's average annual payments per patient are persistently close to the cap. For each hospice program j, we computed $\text{Dist}_j := T_j^{-1} \sum_t |\text{PPP}_{jt}/\text{Cap}_t - 1|$ and $\text{Census}_j := T_j^{-1} \sum_t \text{Census}_{jt}$, where T_j is the number of fiscal years that program j is observed between FY 2001 and FY 2018. Figure A6 plots the joint distribution of $(\text{Dist}_j, \text{Census}_j, T_j)$. It shows that most hospice programs are not persistently close to the cap. It also shows that those hospice programs that are persistently within 10% of the cap are either observed for a short time or have small patient censuses—with one exception: the figure shows that Vitas Healthcare Corporation of Florida operated for 15 fiscal years between FY 2001 and FY 2018 and, in this time, it had an average annual patient census of approximately 11,000—the highest in the country—and an average annual payment per patient within 4.4% of the cap per year on average. Figure A7 plots trends in its $\text{PPP}_{jt}/\text{Cap}_t$. It shows that between 2006-2008, Vitas consolidated its operations in Florida into one program. The program's average annual payments per patient were almost exactly equal to the cap between FY 2009 and FY 2018. In other words, Vitas appears to have nearly maximized average annual payments per patient in Florida for ten straight years.

While it is difficult to determine Vitas's strategy, a recent OIG report suggests that Vitas Healthcare Corporation of Florida may have inappropriately billed Medicare. After auditing a sample claims associated with Vitas Healthcare Corporation of Florida between 2017 and 2019, OIG alleged that "Vitas received at least $140 million in improper Medicare reimbursement for hospice services that did not comply with Medicare requirements." According to OIG, these payments were associated with claims for CHC and GIC that were "not supported" by clinical records. See OIG (2022b).

D Difference-in-differences

D.1 Robustness checks

We examined whether our results were robust to three alternative difference-in-differences specifications. First, we examined whether they are robust to adding 60 days to either side of the event windows.[38] Hospice programs are still placed into $\mathcal{T}_e$ and $\mathcal{U}_e$ based on their running average annual payments per patient at the beginning of the baseline 180-day window. Second, we examined whether our results are robust to excluding the state, ownership, and wage index controls. Finally, we examined whether our results are robust to excluding programs whose running average annual revenue was within 90-100% of the cap at the start of the event windows. These robustness checks will respectively help us assess whether our estimates are sensitive to the event window, control variables, or any potential compensating adjustments made by hospice programs that were on track to be below—but close to—the cap.

Figure A8 plots the estimates computed from these alternative specifications alongside the estimates from our main specifications. Estimates from the "longer window" specification suggest that hospice programs that are on track to exceed the cap in an outgoing fiscal year have differentially lower enrollment rates in the outgoing fiscal year before the outgoing fiscal year's last three months. Overall, these estimates suggest that our main qualitative findings are not sensitive to these alternative specifications.

D.2 Heterogeneity analysis

We explored whether the difference-in-differences in our main analysis meaningfully vary across ownership type, cap census, proximity to other hospice programs on track to exceed the cap, and affiliation with a nursing home. We interacted our main treatment indicators D with the following covariates. In order to examine heterogeneity by profit status, we used an indicator for whether hospice j was for-profit in the outgoing fiscal year. In order to examine heterogeneity across cap census, we used the log of program j's cumulative cap census in at the start of each event window. In order to examine heterogeneity across proximity to other hospice programs on track to exceed the cap, we used the log of program j's proximity to programs on track to exceed the cap in the outgoing fiscal year, as described in appendix B.8. In order to examine heterogeneity by association with a nursing home, we used an indicator for whether hospice j shared any patient-day with a nursing home in the outgoing fiscal year. Figure A9 reports the results. We do not find statistically significant evidence of meaningful heterogeneity.

[38] We therefore incidentally exclude a small number of hospice programs that are not continuously observed throughout the new 300-day event window but were observed continuously throughout the baseline 180-day event window.

	Samples			
	Full	For-profit	Cap census above annual median	Ever over the cap
	(1)	(2)	(3)	(4)
Estimates				
Test statistic	0.0962	0.1788	-0.9371	2.0079
p-value	0.9234	0.8581	0.3487	0.0447
Upper bandwidth				
Range	0.143	0.156	0.157	0.135
Observations	6,602	5,795	1,432	6,602
Effective observations	2,622	2,410	909	2,497
Lower bandwidth				
Range	0.198	0.187	0.222	0.150
Observations	50,810	23,683	27,273	12,449
Effective observations	6,740	5,132	3,550	3,528

Tab. A1. This table reports McCrary test statistics of the null hypothesis that the distribution of average annual payments per patient relative to the cap is continuous at 1. The test statistics support the idea that hospice programs' average annual payments per patient do not bunch around the cap. "Above annual median," means that a program-year's cap census was above that year's median cap census. The number of "effective observations" in each column is the number of observations within the upper and lower bandwidths. The bandwidths are selected based on the mean squared error of the density estimators (Cattaneo et al., 2018). Figure 2 plots this distribution for column (1) (i.e., the full distribution). See the discussion near page 16.

	1[Over cap]	Liability	Liability per unique patient	Gross Medicare payments	PPP
	(1)	(2)	(3)	(4)	(5)
Log(WI)	0.115***	4.544***	2.870**	0.226	0.415***
	(0.023)	(1.047)	(1.138)	(0.169)	(0.063)
Estimator	OLS	PPML	PPML	PPML	PPML
Hospice FE	Y	Y	Y	Y	Y
Period FE	Y	Y	Y	Y	Y
Age & ownership FE	Y	Y	Y	Y	Y
Only ever over	-	-	-	-	-
$\bar{Y}$	0.113	177,361	3,106	3,673,515	15,163
N	56,193	17,550	17,480	56,202	56,193
Clusters	5,246	2,082	2,075	5,246	5,246

	Medicare payments per patient-day			
	Gross	Net	Gross	Net
	(6)	(7)	(8)	(9)
Log(WI)	0.328***	0.301***	0.341***	0.203***
	(0.038)	(0.037)	(0.060)	(0.056)
Estimator	PPML	PPML	PPML	PPML
Hospice FE	Y	Y	Y	Y
Period FE	Y	Y	Y	Y
Age & ownership FE	Y	Y	Y	Y
Only ever over	-	-	Y	Y
$\bar{Y}$	147	143	147	136
N	56,202	56,202	18,425	18,425
Clusters	5,246	5,246	2,233	2,233

	Annual cap census contribution per patient	Cap census	Number of unique patients	Number of patient-days
	(10)	(11)	(12)	(13)
Log(WI)	0.002	-0.103	-0.084	-0.044
	(0.009)	(0.083)	(0.088)	(0.125)
Estimator	PPML	PPML	PPML	PPML
Hospice FE	Y	Y	Y	Y
Period FE	Y	Y	Y	Y
Age & ownership FE	Y	Y	Y	Y
Only ever over	-	-	-	-
$\bar{Y}$	0.730	273	357	23,619
N	56,202	56,202	56,202	56,202
Clusters	5,246	5,246	5,246	5,246

Tab. A2. This table reports estimates from models relating outcomes to the wage index. Program-level cluster-robust SEs in parentheses. See the discussion in section 5.2.

	1[Over cap]	Liability	Liability per unique patient	Gross Medicare payments	PPP	1[CCN terminated]
	(1)	(2)	(3)	(4)	(5)	(6)
$1[\text{Over cap}]_{jt-1}$	0.193***	1.242***	0.725***	0.074	0.162***	0.056***
	(0.010)	(0.058)	(0.064)	(0.055)	(0.008)	(0.003)
Estimator	OLS	PPML	PPML	PPML	PPML	OLS
Hospice FE	Y	Y	Y	Y	Y	-
Period FE	Y	Y	Y	Y	Y	Y
Age & ownership FE	Y	Y	Y	Y	Y	Y
$\bar{Y}$	0.112	191,840	2,985	3,911,340	15,261	0.025
N	51,203	15,497	15,451	51,210	51,203	52,918
Clusters	4,878	1,893	1,887	4,879	4,878	5,375

	Medicare payments per patient-day		Annual cap census contribution per patient	Cap census	Number of unique patients	Number of patient-days
	Gross	Net				
	(7)	(8)	(9)	(10)	(11)	(12)
$1[\text{Over cap}]_{jt-1}$	-0.008***	-0.063***	-0.074***	-0.107**	-0.025	0.095**
	(0.002)	(0.004)	(0.003)	(0.046)	(0.044)	(0.044)
Estimator	PPML	PPML	PPML	PPML	PPML	PPML
Hospice FE	Y	Y	Y	Y	Y	Y
Period FE	Y	Y	Y	Y	Y	Y
Age & ownership FE	Y	Y	Y	Y	Y	Y
$\bar{Y}$	147	144	0.726	286	376	25,053
N	51,210	51,210	51,210	51,210	51,210	51,210
Clusters	4,879	4,879	4,879	4,879	4,879	4,879

Tab. A3. This table reports estimates from models relating outcomes to a 1-year lag indicator of whether a hospice program exceeded the cap. Program-level cluster-robust SEs in parentheses. See the discussion in section 5.3.

		Above-cap		Below-cap		
	All	All	W/I 10%		≥80%	All
	(1)	(2)	(3)	(4)	(5)	(6)
Demographic characteristics						
1[Female]	0.57	0.59	0.59	0.58	0.58	0.57
1[White]	0.88	0.80	0.82	0.82	0.84	0.88
1[Black]	0.08	0.13	0.10	0.10	0.09	0.08
1[Asian]	0.01	0.01	0.01	0.01	0.01	0.01
1[Hispanic]	0.02	0.04	0.04	0.05	0.03	0.02
1[N. Amer. Native]	0.00	0.00	0.00	0.00	0.00	0.00
Age at enrollment	81.7	81.7	82.3	82.4	82.3	81.7
Lifetime hospice utilization						
Hospices	1.06	1.22	1.15	1.11	1.10	1.05
1[Ever multiple hospices]	0.05	0.18	0.13	0.10	0.09	0.05
Live discharges	0.22	0.60	0.44	0.34	0.32	0.20
1[Ever live discharged]	0.17	0.39	0.31	0.25	0.24	0.16
Lifetime LOS in hospice care	90	200	172	148	142	85
Characteristics pertaining to mortality						
1[Decedent]	0.98	0.95	0.96	0.97	0.97	0.99
Age at death (decedents)	82.1	82.9	83.1	83.1	82.9	82.1
Remaining days of life (decedents)	133	348	256	210	201	123
Remaining days of life	156	425	308	248	236	143
1[Died in hospice]	0.92	0.81	0.86	0.89	0.90	0.92
Consec. hospice days at EOL	60	98	98	91	88	58
N (patients, M)	15.6	0.7	0.3	0.7	1.7	14.9

Tab. A4. This table reports descriptive statistics about the patient-level data. On average, patients enrolling in program-years that exceed the cap have longer average lifetime LOS in hospice care, are treated at more hospice programs, and experience more live discharges. They have longer remaining lifetimes and are less likely to ultimately die in hospice care. Each patient is linked to a program-year based on the program and fiscal year associated with their first day in hospice care. They are assigned to a column based on the program's cap liability in that day's fiscal year. For non-decedents, remaining days of life is calculated through the end of our sample period (i.e., 12/31/2019).

		Above-cap		Below-cap		
	All	All	W/I 10%		≥80%	All
	(1)	(2)	(3)	(4)	(5)	(6)
Age and ownership						
Age	12.27	6.20	7.28	8.27	8.76	13.06
1[For profit]	0.52	0.88	0.87	0.84	0.83	0.47
1[Not for profit]	0.36	0.06	0.07	0.09	0.10	0.40
1[Government]	0.05	0.01	0.01	0.01	0.01	0.06
Enrollments and live discharges						
New enrollments	271	106	158	248	242	293
Live discharges	56	57	62	77	72	56
% hospitalized within 7 days of enrollment	0.34	0.23	0.26	0.28	0.29	0.36
% in a NH within 7 days of enrollment	0.27	0.19	0.24	0.28	0.29	0.28
% elsewhere within 7 days of enrollment	0.47	0.63	0.57	0.52	0.50	0.45
Staffing						
Skilled nurse visits per patient-day	0.23	0.19	0.21	0.21	0.21	0.23
Social worker visits per patient-day	0.05	0.03	0.04	0.04	0.04	0.05
Home health aide visits per patient-day	0.24	0.25	0.27	0.28	0.28	0.24
Levels of hospice care						
Fraction of patient-days with RHC	0.98	0.99	0.99	0.99	0.99	0.98
Fraction of patient-days with CHC	0.00	0.00	0.00	0.00	0.00	0.00
Fraction of patient-days with IRC	0.00	0.00	0.00	0.00	0.00	0.00
Fraction of patient-days with GIC	0.01	0.00	0.00	0.00	0.01	0.01
Other characteristics						
1[Patient-day at a NH]	0.22	0.13	0.18	0.22	0.24	0.24
1[At least one NH patient]	0.96	0.90	0.94	0.96	0.96	0.96
Proximity to other over-cap programs	0.11	0.24	0.20	0.18	0.17	0.10
1[CCN terminated in next FY]	0.03	0.08	0.04	0.02	0.02	0.02
N (program-years)	57,412	6,602	2,018	2,926	6,824	50,810
Unique programs	5,685	2,214	1,279	1,605	2,517	5,300

Tab. A5. This table reports descriptive statistics about the (program, fiscal year)-level data. On average, programs that exceed the cap are newer, more likely to operate for profit, and have a higher proximity to other programs that exceed the cap contemporaneously. They enroll fewer patients, live discharge a larger fraction of patients, and enroll fewer patients with a recent hospital or nursing home stay. Finally, they are more likely to close in the following year.

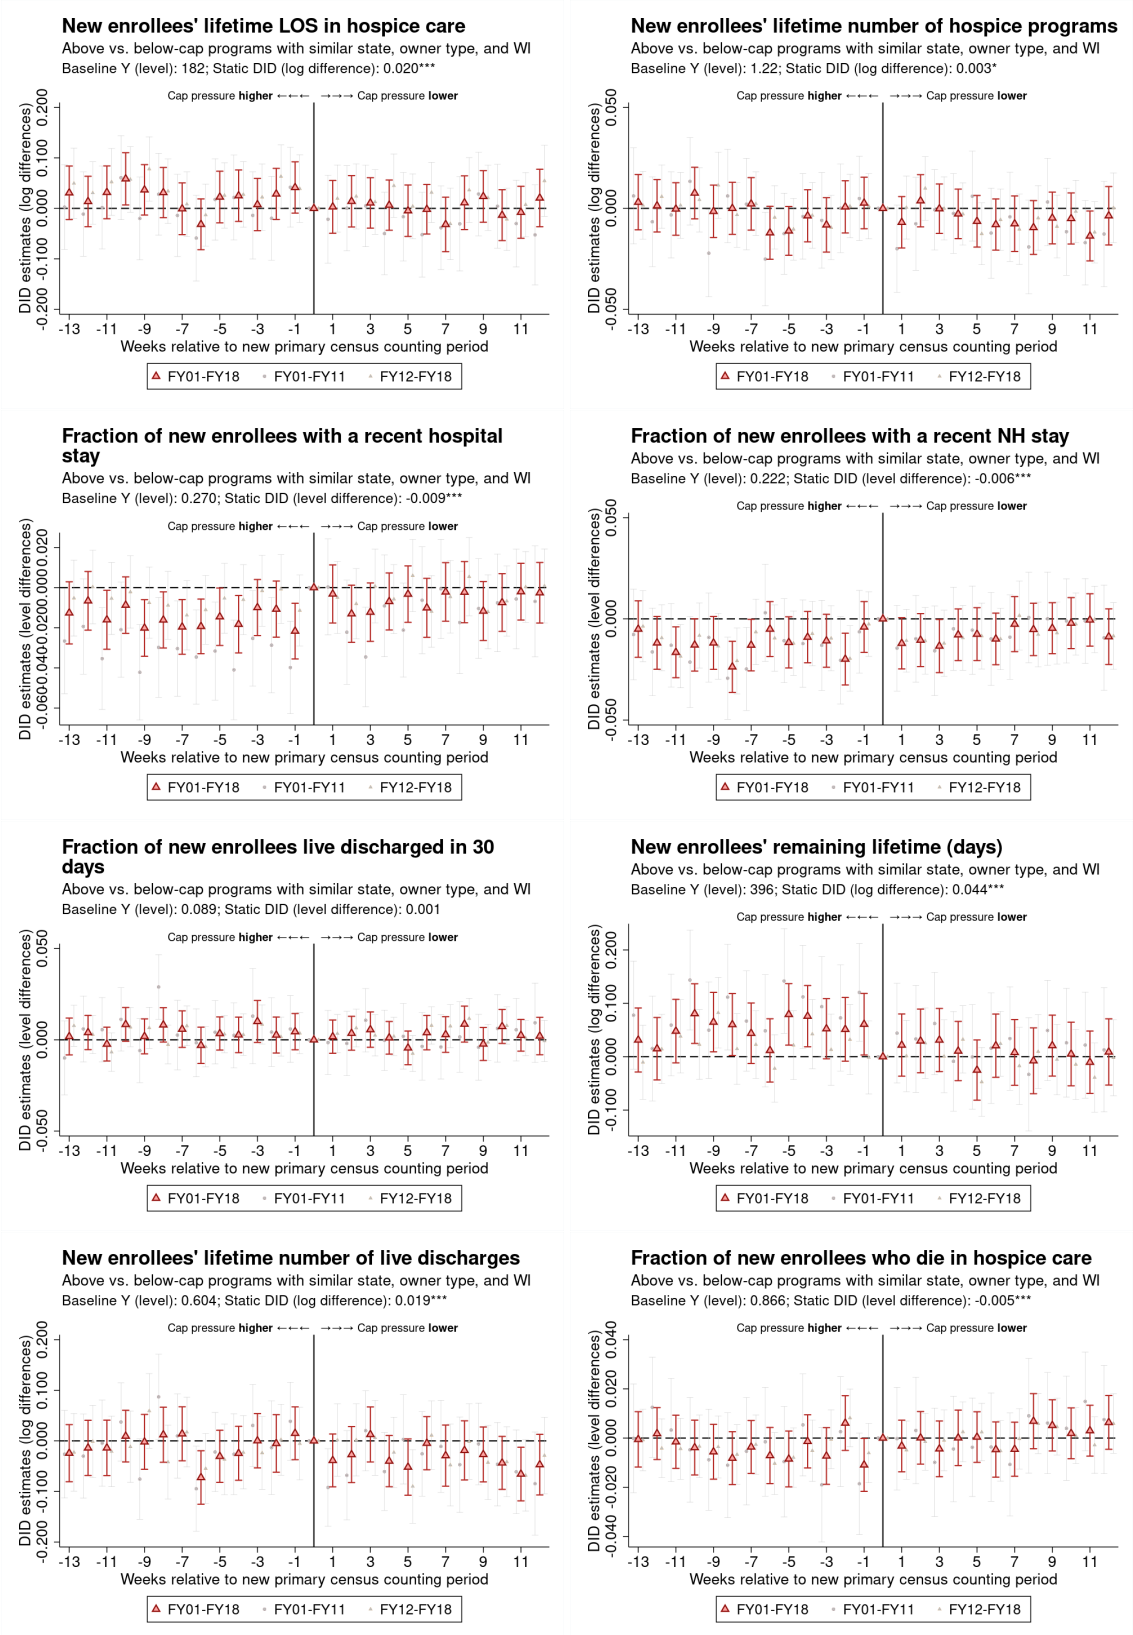


Fig. A1. This figure plots dynamic estimates for new enrollees in our DID analysis. The 95% confidence intervals are computed with program-level cluster-robust SEs. The heading reports statistics estimated from the full sample period. Baseline Y is the average value of the outcome during relative time 0 among programs in $\mathcal{T}$. Static DID is an estimate of β from equation (5). Observations are weighted by the number of enrollees. The primary census counting periods are described in appendix A.

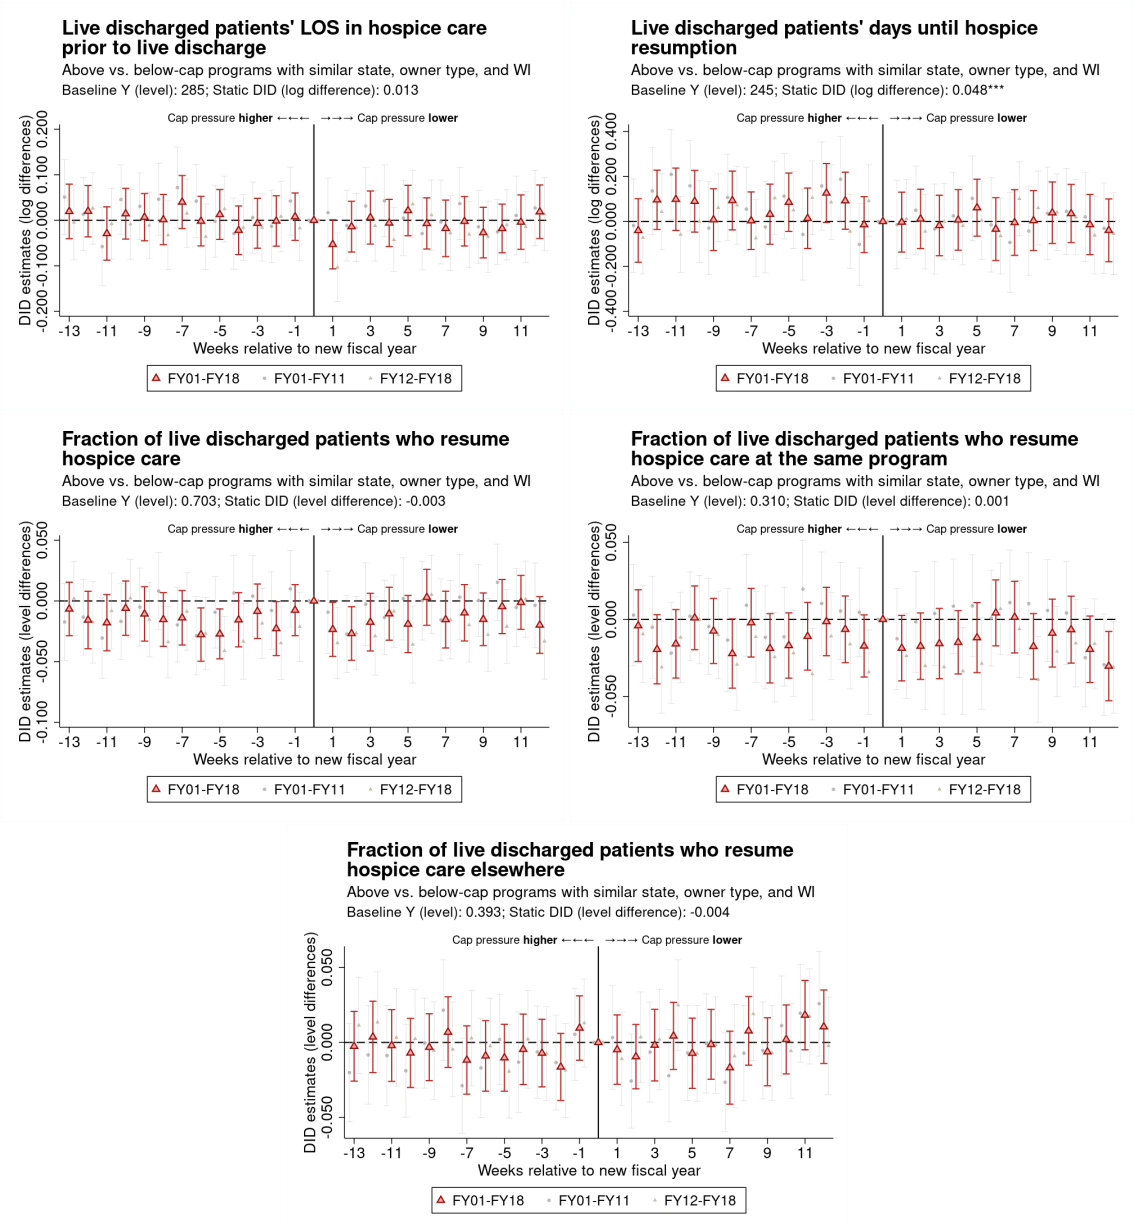


Fig. A2. This figure plots dynamic estimates for live discharged patients in our DID analysis. The 95% confidence intervals are computed with program-level cluster-robust SEs. The heading reports statistics estimated from the full sample period. Baseline Y is the average value of the outcome during relative time 0 among programs in $\mathcal{T}$. Static DID is an estimate of β from equation (5). Observations are weighted by the number of live discharged patients.

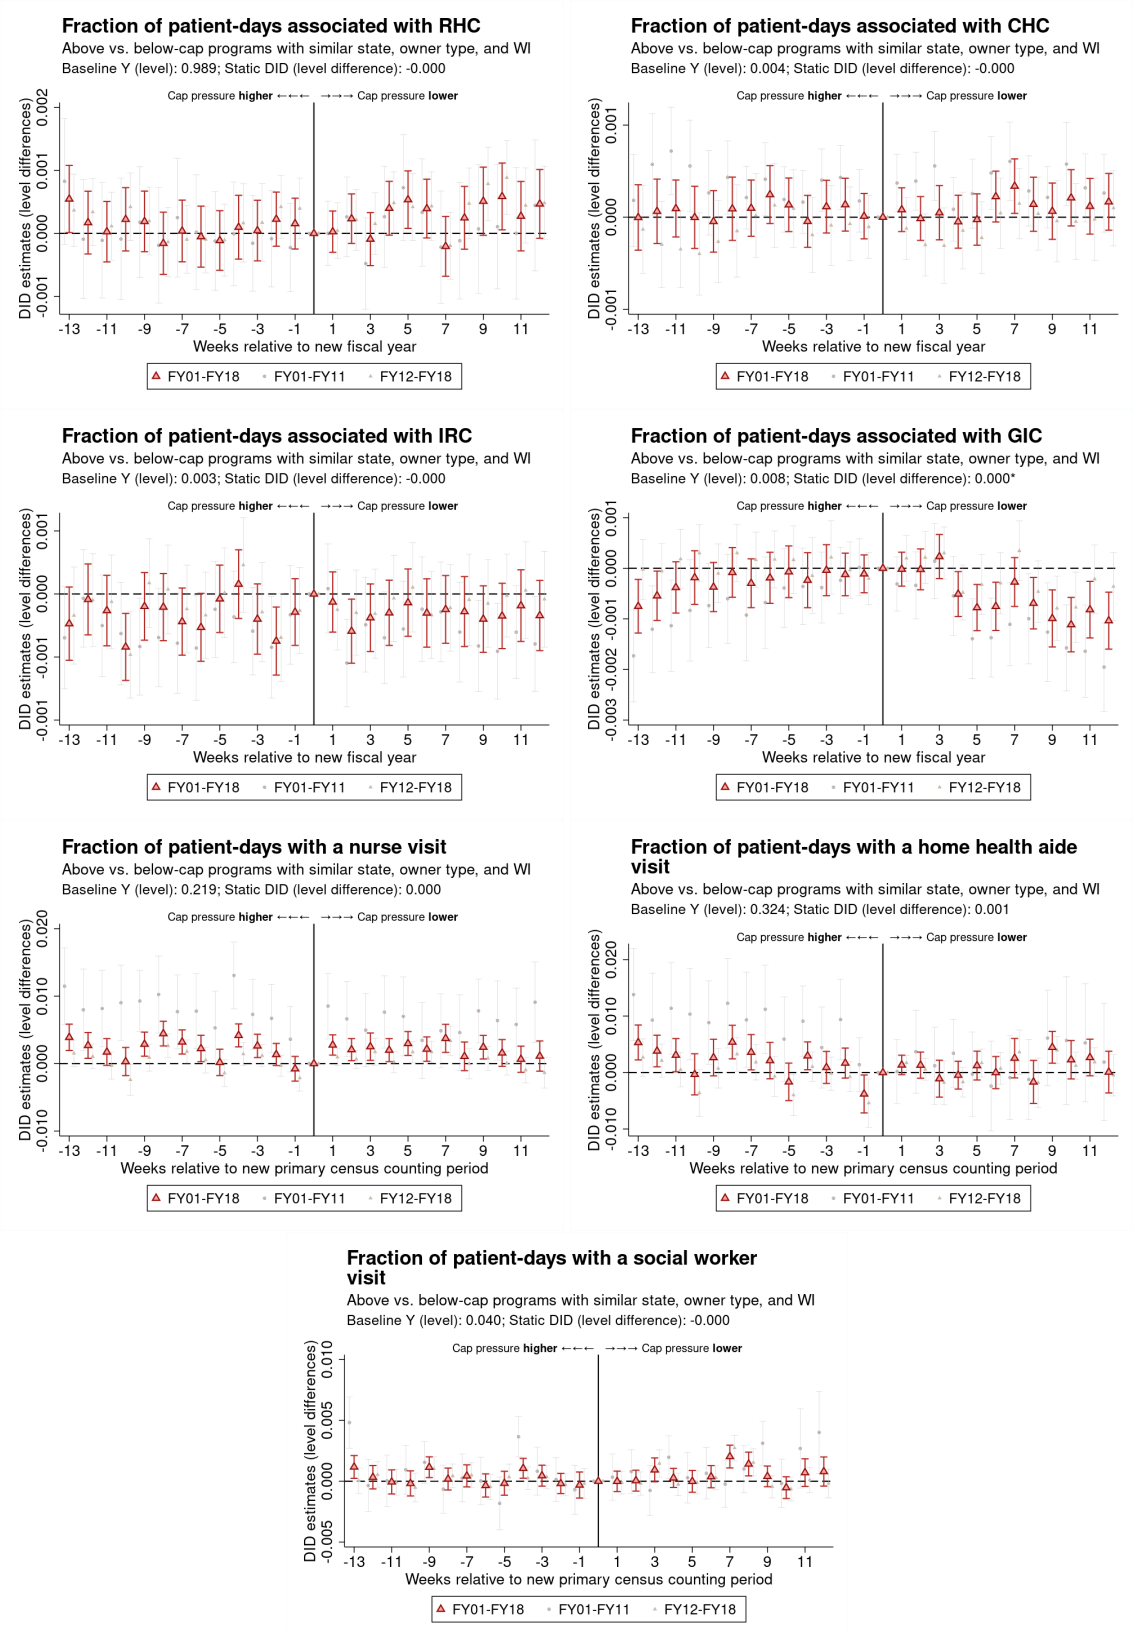


Fig. A3. This figure plots dynamic estimates for active patients in our DID analysis. The 95% confidence intervals are computed with program-level cluster-robust SEs. The heading reports statistics estimated from the full sample period. Baseline Y is the average value of the outcome during relative time 0 among programs in $\mathcal{T}$. Static DID is an estimate of β from equation (5). Observations are weighted by the number of active patients.

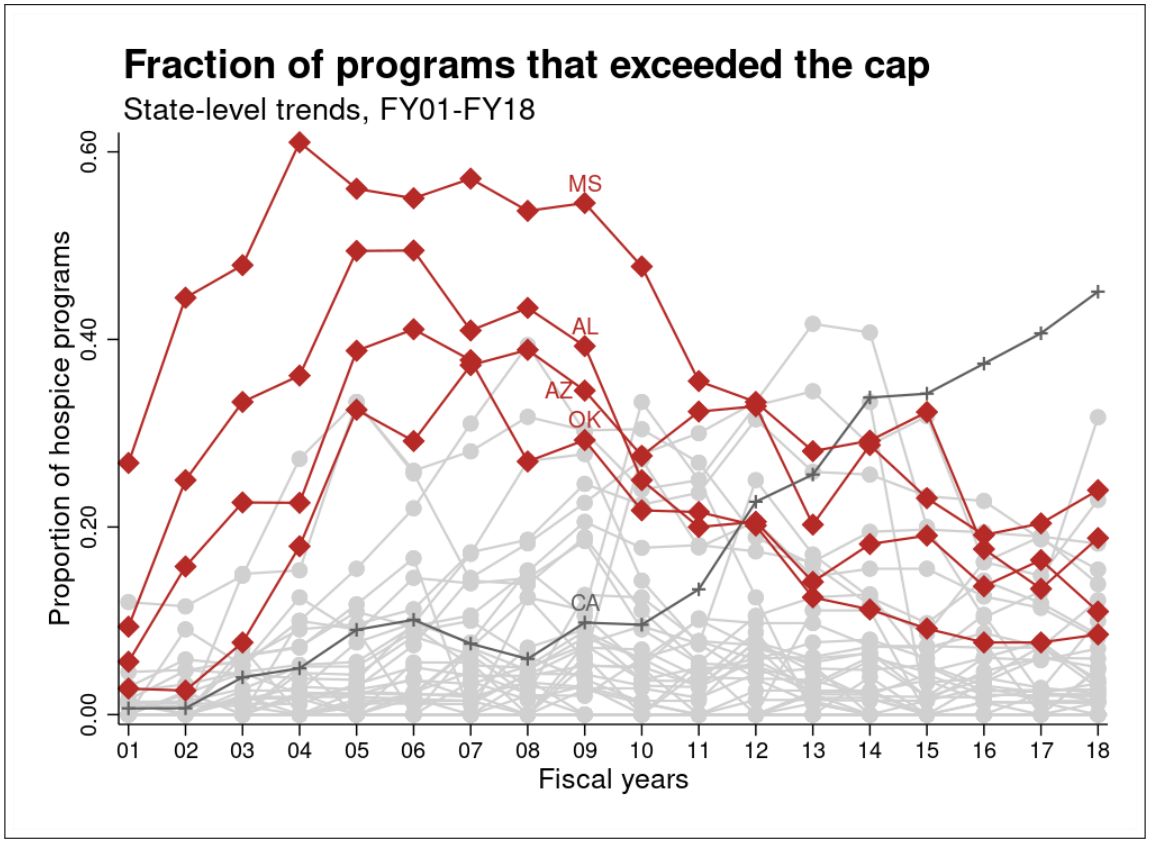


Fig. A4. This figure plots state-level trends in the proportion of hospice programs that exceeded the cap between fiscal years 2001 and 2018. It shows that cap liabilities were particularly prevalent in MS, AL, OK, AZ, and CA during this period. See the discussions in appendix section C.

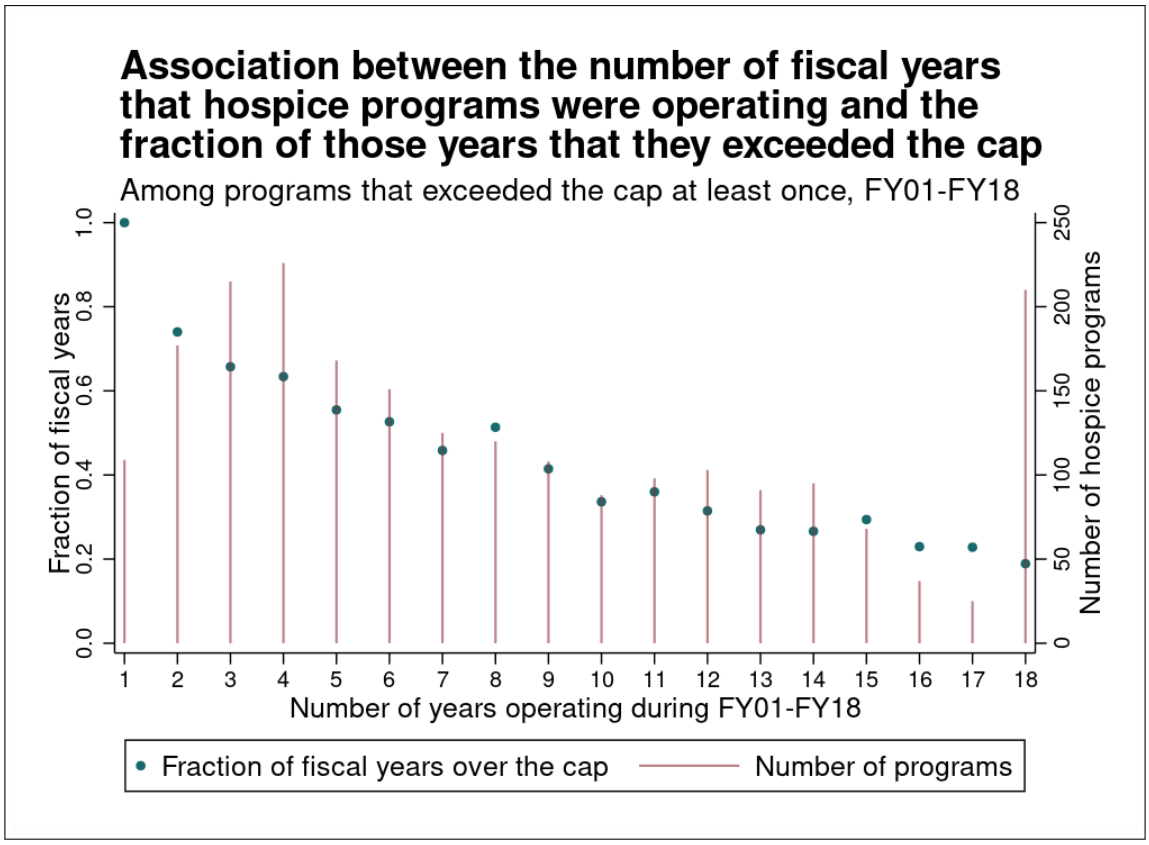


Fig. A5. This figure plots the joint distribution between the number of years that hospice programs operated between FY 2001 and FY 2018 and the fraction of those years that they exceeded the cap. It only includes programs that exceeded the cap at least once. It shows that hospice programs that exceed the cap at least once often exceed it multiple times. See the discussions in appendix section C.

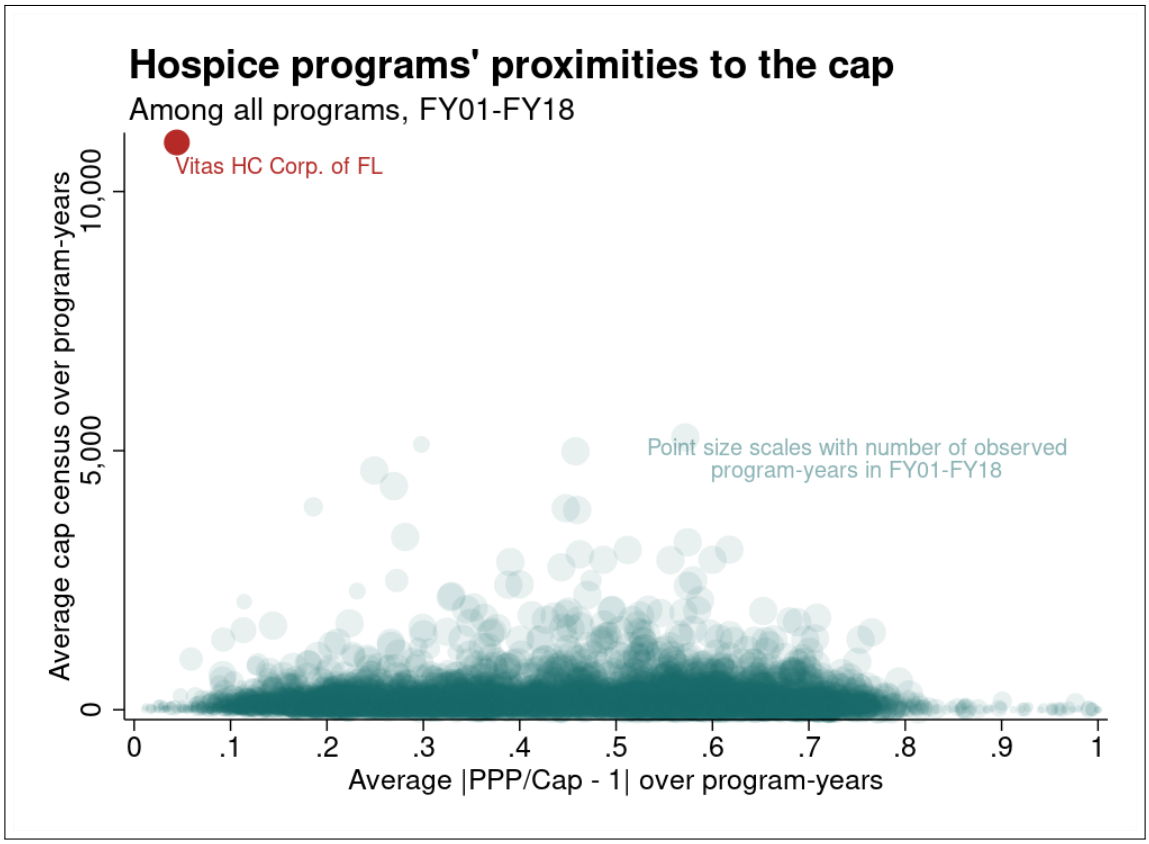


Fig. A6. This figure plots the joint distribution between the number of years that hospice programs are observed in our sample between FY 2001 and FY 2018, their average annual cap censuses, and their average proximities to the cap during this time. It supports the idea that most hospice programs do not persistently equilibrate their average annual payments per patient with the cap. However, it also shows that Vitas Healthcare Corporation of Florida is by far the largest hospice program in the country and persistently close to the cap. The x-axis is censored at 1 for visual clarity. See the discussions in appendix section C.

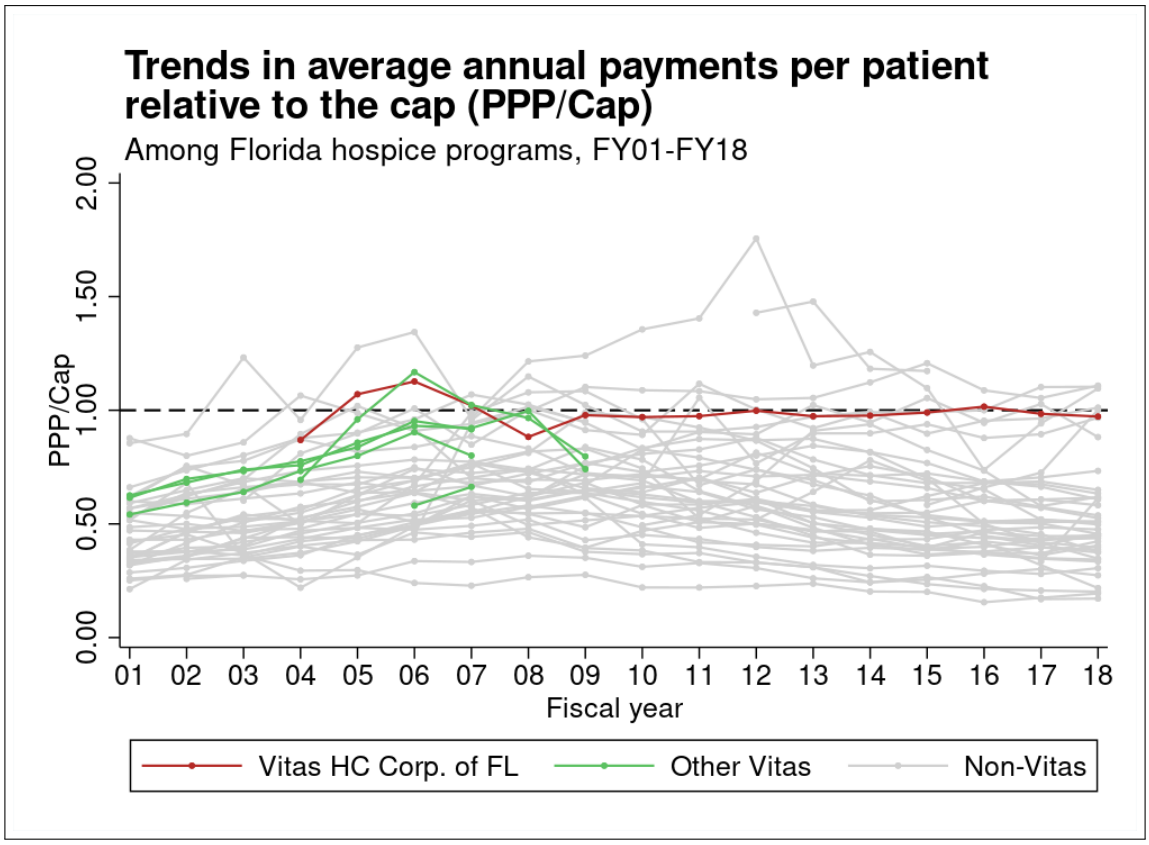


Fig. A7. This figure plots trends in $\text{PPP}_{jt}/\text{Cap}_{jt}$—i.e., hospice programs' average annual payments per patient relative to the cap—among Florida hospice programs. It shows that Vitas Healthcare Corporation of Florida has nearly equilibrated its average annual payments per patient with the cap since 2009. Other Vitas hospice programs were identified by the word "Vitas" in their facility name. See the discussions in appendix section C.

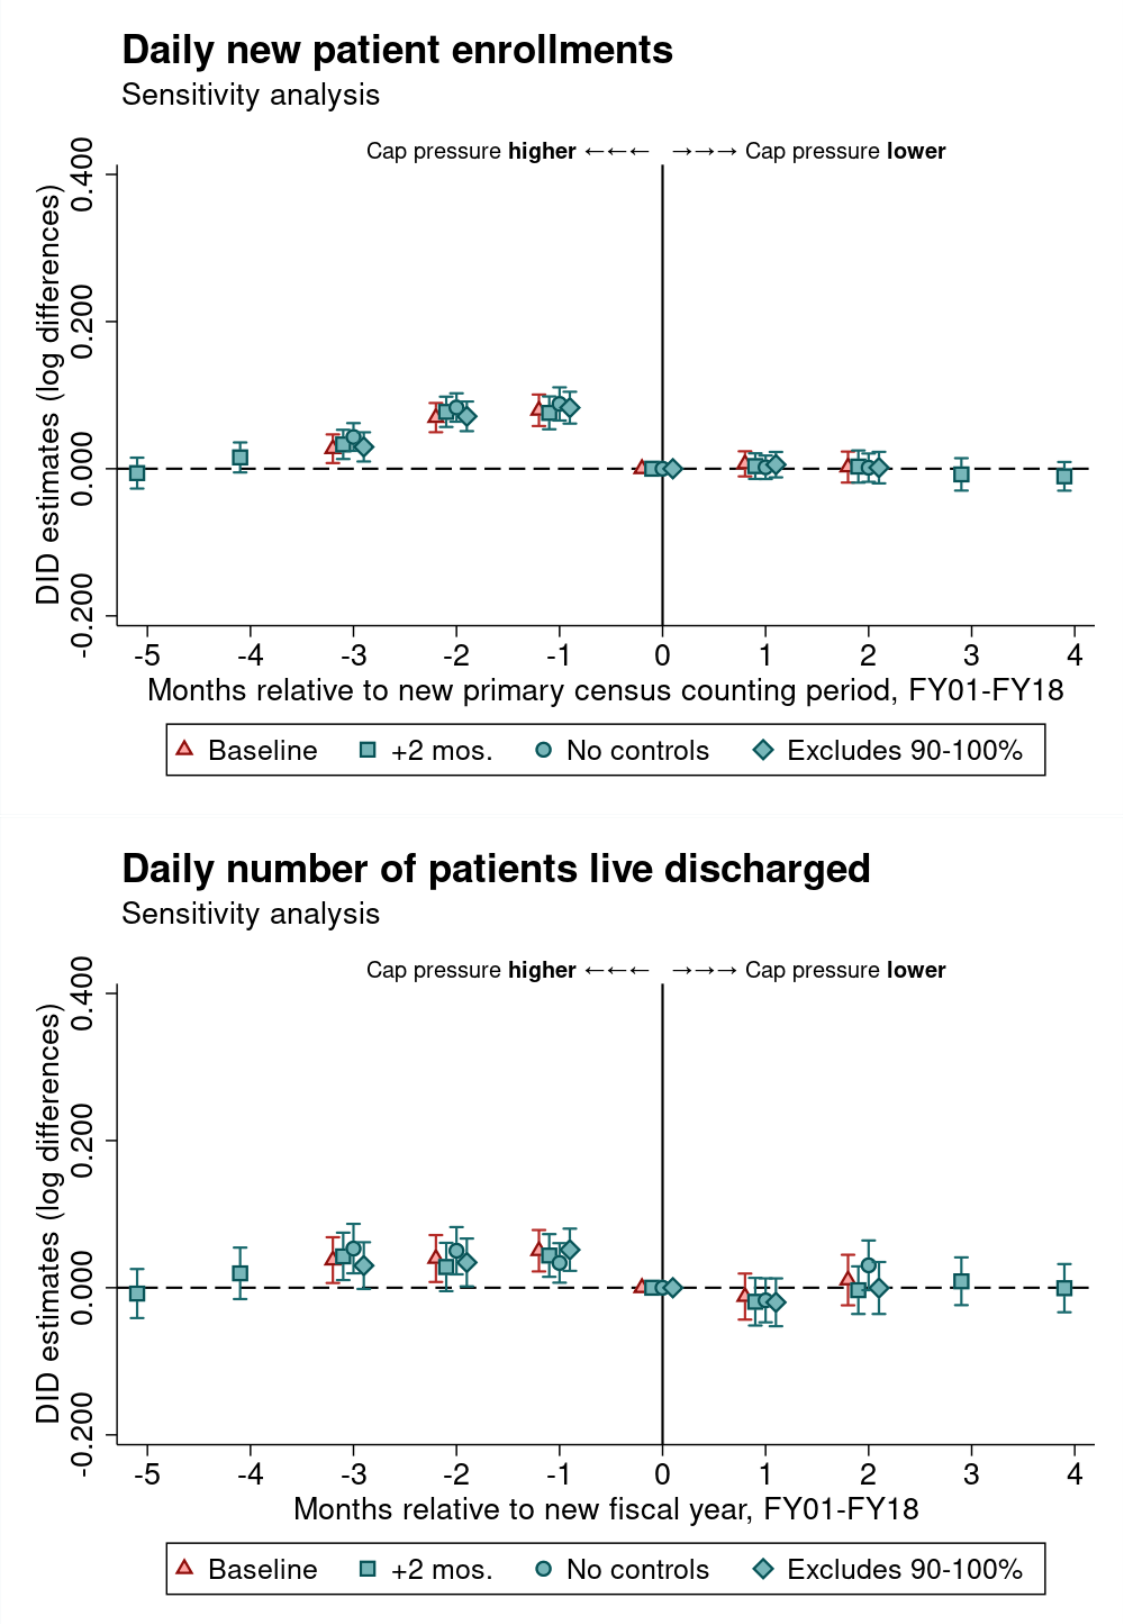


Fig. A8. This figure plots results of the sensitivity analysis discussed in section D.1. It suggests that our baseline estimates are not sensitive to several variations of the DID analysis. The 95% confidence intervals are computed with program-level cluster-robust SEs.

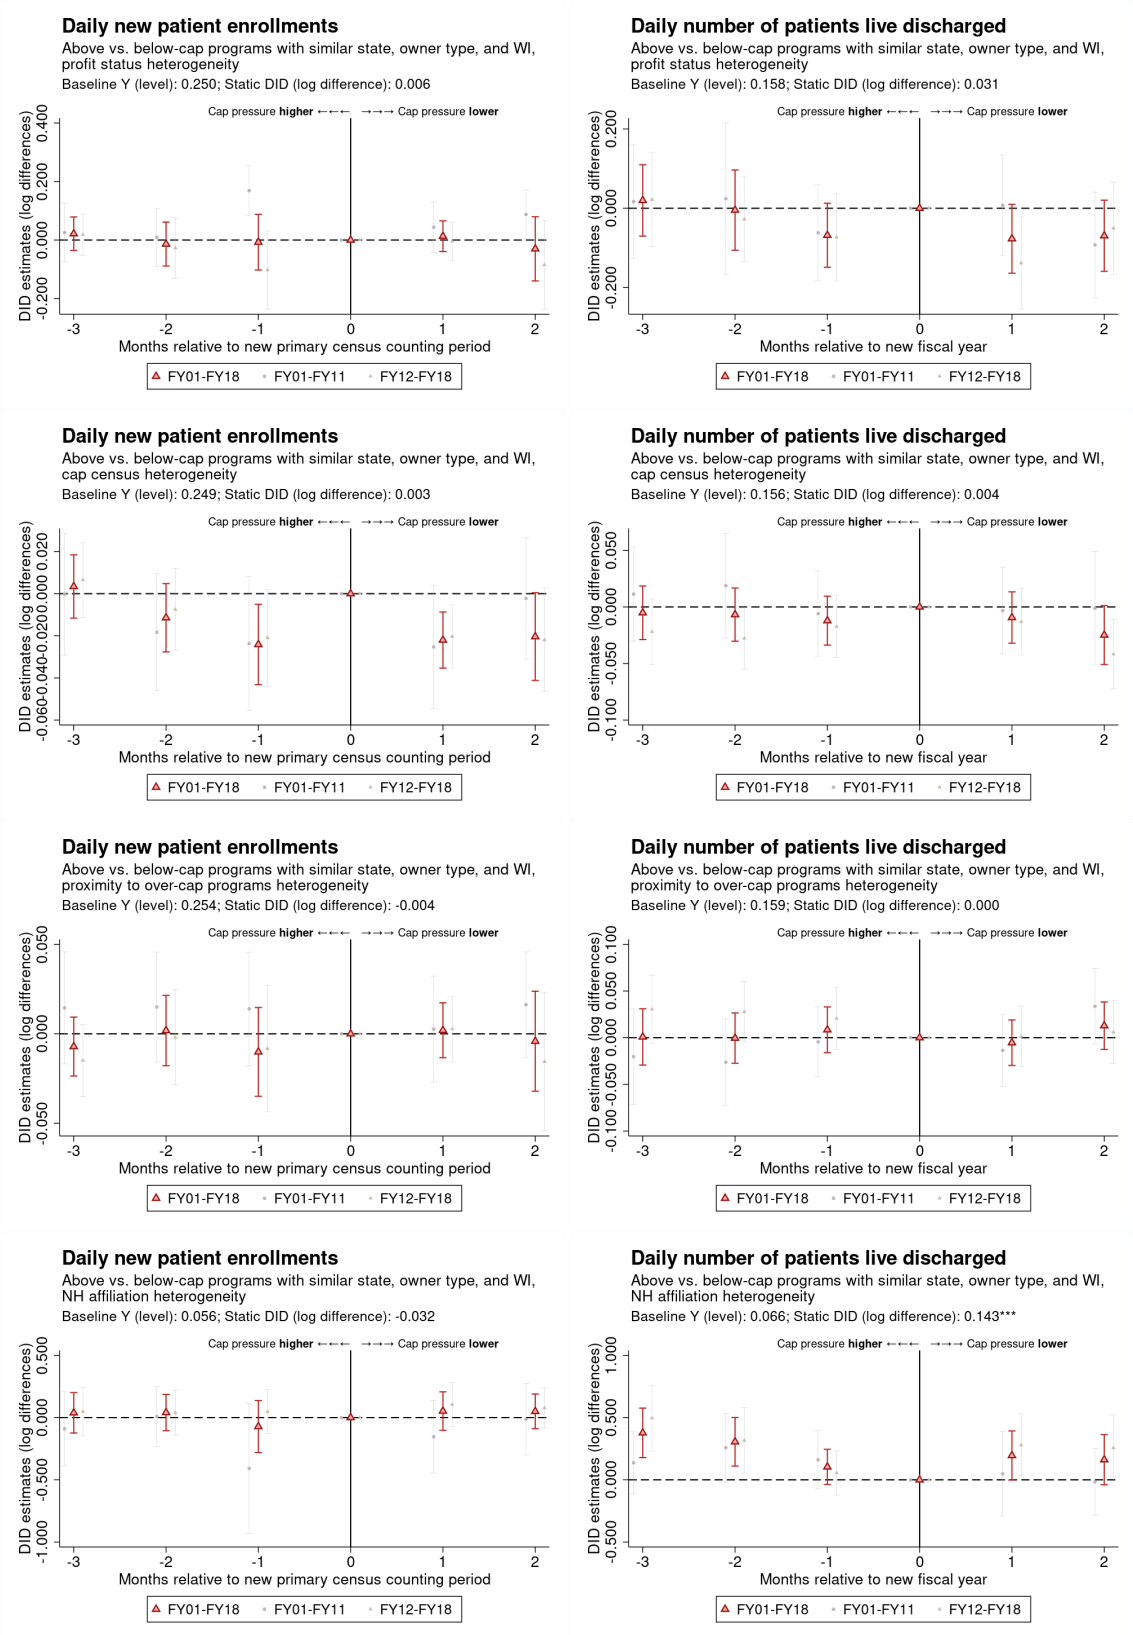


Fig. A9. This figure plots the results of the heterogeneity analysis described in appendix D.2. The 95% confidence intervals are computed with program-level cluster-robust SEs.