PRELIMINARY DRAFT
PLEASE DON'T CIRCULATE OR CITE BEFORE CHECKING WITH ONE OF US

Provider Incentives and Health Care Costs: Evidence from Long-Term Care Hospitals*

Liran Einav, Amy Finkelstein, and Neale Mahoney†

September 2016

Abstract. We study the design of provider incentives in the post-acute care setting – a high-stakes but under-studied segment of the healthcare system – by examining the impact of a sharp jump in payments to long-term care hospitals (LTCHs) that occurs when a patient's stay reaches a pre-specified number of days. The descriptive evidence indicates that discharge decisions from the LTCH respond strongly to approximately $13,000 payment increase for keeping a patient an additional day at the threshold. The marginal patient discharged after the threshold is a relatively healthy one. Despite the large incentives and behavioral response in a high mortality population (90 day mortality is 30%), we are unable to detect any compelling evidence of an impact of the incentives on patient mortality. To quantify the impact of the financial incentives and assess behavior under counterfactual payment schedules, we specify and estimate a simple dynamic discrete choice model of LTCH discharge decisions. We find that our least generous payment schedule could save about $12,000 per patient or $1.5 billion annually. Interestingly, we find that our most generous payment schedule, which provides weakly higher payments at every length of stay, also reduces government spending through its effect on discharge behavior. This result highlights how improved financial incentives can reduce healthcare spending, without negative consequences for patient health and perhaps even industry profits.

JEL classification numbers: D22, I11, L21
Keywords: Healthcare, post-acute care, financial incentives, nonlinear contracts

---

*We are grateful to Yunan Ji, Abby Ostriker, and Yin Wei Soon for excellent research assistance, and to Jeremy Kahn and Hannah Wunsch for helpful comments. Einav and Finkelstein gratefully acknowledge support from the NIA (R01 AG032449).

†Einav: Department of Economics, Stanford University, and NBER, leinav@stanford.edu; Finkelstein: Department of Economics, MIT, and NBER, afink@mit.edu; Mahoney: Booth School of Business, University of Chicago, and NBER, Neale.Mahoney@chicagobooth.edu.

1 Introduction

Healthcare spending is one of the largest fiscal challenges facing the U.S. federal government. In 2014, the U.S. federal government spent $1.1 trillion on public healthcare programs (BEA, 2015) and the CBO projects that spending will grow to $2.0 trillion by 2026 (CBO, 2016). Within the healthcare system, post acute care (PAC) is an important but relatively under-studied sector. Post-acute care is the term used for care provided to help patients recover from a surgery or other acute care event, and is defined as care provided by a long-term care hospital (LTCH), skilled nursing facility (SNF), inpatient rehabilitation facility (IRF), or home health care agency (HHA) (MedPAC, 2015b). In the Traditional Medicare (TM) program that we study, PAC spending was $60 billion or about 16 percent of total program spending in 2013, and is growing at a faster rate than overall Medicare spending (MedPAC, 2004, 2015a).

The nature of the post acute care system makes it ripe for inefficiencies. Patients take complex paths through PAC, often receiving services from two or more facilities during a single episode of care. The functional distinction between different facilities is not always clear, allowing providers to exert substantial discretion over the location of treatment and length of stay. A recent Institute of Medicine report found that – despite accounting for only 16 percent of spending – PAC contributed to a striking 73% of the unexplained geographic variation in Medicare spending, underlining concerns about provider discretion and inefficiencies in the sector (Newhouse et al., 2013).

In this paper, we study the role of financial incentives in determining patient flows and government spending in the Medicare PAC system. Given its fiscal importance, understanding the effects of financial incentives is a natural area for inquiry. Moreover, inefficiencies in the PAC sector have potentially important implications for public health. Over 40% of hospital patients are discharged to PAC (MedPAC, 2015b). These PAC stays are disproportionately concentrated in high-risk patients who might be more vulnerable to inefficiencies in the delivery of care.[1] However, despite the importance of the PAC system, it has received relatively little attention from academic economists.

Our analysis focuses on patients whose point of entry into the PAC system is an LTCH. We focus on these patients because of sharp variation in provider incentives at this type of facility. This is illustrated in Figure 1: providers are reimbursed a daily amount (of approximately $1,300) up to a threshold number of days, at which point there is a large (approximately $13,000) lump sum payment for keeping a patient an additional day beyond the threshold, but no payments for any additional days beyond it. We investigate the effects of this "jump" in payments using detailed Medicare claims and administrative data on the universe of LTCH stays over the 2008-2012 period – when this non-linear payment schedule was in effect – as well as a 2000-2002 period, when LTCHs were instead reimbursed under a linear payment schedule.

We start by presenting descriptive evidence on the effect of the jump in payments. Discharges respond strongly to the payment increase, with the share of stays discharged increasing from 2% to 9% at precisely the date of the payment increase. The marginal patient discharged at the threshold appears to be (relatively) much healthier than average. At the threshold, patients are

[1] We calculate that 15% of Medicare deaths involve a PAC stay in the 30 days prior to death.

disproportionately likely to be discharged to "downstream" (to a less intensive PAC facility or home) relative to "upstream" (e.g. to acute care hospitals) locations. And at the threshold, the 30-day mortality rate of discharged patients is substantially lower than patients who are discharged earlier.

A natural question raised by this evidence is whether distortions in the timing of discharge – and therefore distortions to patients' location of treatment – has an impact on patient health. Given the high mortality rate for LTCH patients (16% die within 30 days of LTCH admission), if the distortions are harmful, it seems plausible that we could detect an impact on mortality in our setting. However, empirical analysis of mortality effects is challenging because, unlike discharge behavior, they are not expected to be visually present right "at" the threshold.

The available evidence shows no compelling evidence of any mortality effects from the distortions in discharge behavior. We find no evidence of a change in the level or the slope of the mortality hazard in the vicinity of the jump in payments. An additional informative contrast is provided by comparing the experience of for-profit and non-profit LTCHs. We show that while both types of hospitals have similar payment schedules, the behavioral response in discharges is much stronger for-profit hospitals. Despite the larger behavioral response of for-profit LTCHs, the two types of facilities have virtually identical mortality patterns, both in the immediate vicinity of the threshold – suggesting no on-impact effect of the location distortion on mortality – and also in the weeks before and after the threshold – indicating that the distortion in discharge behavior does not have a less-immediate, gradual effect on mortality. These results suggest that the marginal patient is able to receive similar care whether they are located in an LTCH or in their alternative setting, which empirically is usually a SNF.

The descriptive analysis provides compelling evidence that providers respond to financial incentives, but does not provide a natural way for gauging the magnitude of this response or estimating how treatment patterns and Medicare costs would be effected by counterfactual payment schedules. To address these questions, in the second half of the paper we specify and estimate a stylized dynamic model of LTCH discharge behavior. The LTCH faces a (daily) decision of whether to retain the patient (characterized by their health, which evolves stochastically over time) or discharge them to another facility. The provider's objective function includes both net revenue (Medicare payments net of costs) and patient utility. If the patient is discharged from the LTCH, the provider receives zero net revenue, but internalize at least a portion of the patient utility from being treated in an alternative location. If the LTCH keeps the patient, it receives a net revenue that depend on Medicare's payment schedule, while also accounting for the patient's utility from being treated in the LTCH and the option value of making a similar discharge decision the following day. The provider problem can therefore be described by a standard dynamic discrete choice problem.

We estimate the model by simulated method of moments to match the observed discharge and mortality patterns under the linear and non-linear payment schedules. To take advantage of the variation provided by the sharp jump in payments, we assign greater weights to moments that are close to the payment jump. The estimated model fits the data reasonably well. We use our model and the estimated parameters to investigate the effects of two types of counterfactual exercises. In our primary exercise, we consider three alternative payment schedules, which all eliminate the "jump" in LTCH payments at the threshold, but in different ways. We estimate that the least generous reimbursement schedule could save about $12,000 per patient, or about $1.5 billion annually. Interestingly, even the most generous reimbursement schedule, which from an accounting perspective is even more generous that the observed schedule, is estimated to lead to a (small) reduction in Medicare spending through its effect on discharge behavior. In our second set of counterfactuals, we return to the variation in the response to incentives between for-profit and non-profit LTCHs, and estimate that if all for-profit LTCHs behaved like non-profits, government spending would be lower by approximately $2,000 per patient, or about $250 million per year.

Given the importance of healthcare spending in the economy and in public sector budgets, it is not surprising that there exists a large literature examining how healthcare spending responds to financial incentives. What is surprising – and arguably unfortunate from this perspective – is that the vast majority of this literature (including much of our own prior work) has focused on the impact of *consumer* financial incentives (such as deductibles and co-payments).[2] The majority of healthcare spending, however, is accounted for by a small share of high-cost individuals whose spending is largely in the "catastrophic" range where deductibles and co-insurance often no longer bind, and thus where consumer cost-sharing is likely to have little impact relative to provider-side incentives and reimbursement policies.[3]

The relative lack of research on the provider side presumably reflects the difficulties in finding clean variation in incentives to model and study. Perhaps not surprisingly therefore, the sharp incentives created by the current LTCH payment schedule have already received some attention in both academic (Kim et al., 2015) and popular (Weaver et al., 2015) spheres. Our descriptive work on discharges around the threshold is quite similar to this prior work, while our analysis of the health of the marginal dischargee and our exploration of mortality effects is new. Our paper is most closely related to Eliason et al. (2016) who – in independent ongoing work – also study the impact of the LTCH payment schedule on discharge behavior descriptively and through the lens of a dynamic model.

Finally, from a more conceptual perspective, our paper is related to a growing literature that seeks to interpret descriptive evidence of the behavioral responses to non-linear payment schedules ("bunching") through the lens of richer economic models that allow for assessments of behavior under counterfactual schedules (Chetty et al., 2011; Einav et al., 2015, 2016; Manoli and Weber,

[2]The literature on the impact of consumer incentives ("moral hazard") in health insurance is too vast to try to summarize or cite here. Most of the work on provider-side responses has focused on descriptive evidence that providers do respond to incentives, with much of the evidence coming from the response to the introduction of the Inpatient Prospective Payment System in 1983 (Cutler and Zeckhauser (2000) provide a review of this evidence). More recently, Clemens and Gottlieb (2014) provide a rare look at the behavioral response of physicians to financial incentives.

[3]For instance, using data from 2013 in the Medical Expenditure Panel Survey, individuals in the top 10% of the spending distribution accounted for 63% of total spending, but had an out-of-pocket share of 6.9% (relative to a 12% out-of-pocket share for the entire population). The top 1% of the distribution accounted for 22% of total spending and had an out-of-pocket share of 4.2%.

forthcoming).

The rest of the paper proceeds as follows. Section 2 provides some background on the PAC sector in the U.S., LTCHs, and our data. In Section 3, we describe the discharge and mortality patterns around the jump in payments. Section 4 presents the model, discusses estimation, and presents the results from our counterfactuals. Section 5 concludes.

2 Setting and Data

2.1 Post-Acute Care in The United States

Post-acute care (PAC) is the term for rehabilitation and palliative services provided to patients recovering from an acute care hospital stay. In the United States, the Center for Medicaid and Medicare Services (CMS) associates PAC with four types of facilities: long-term care hospitals (LTCHs), skilled nursing facilities (SNFs), inpatient rehabilitation facilities (IRFs), and home health agencies (HHAs) (MedPAC, 2015b). In 2013, Medicare paid $60 billion to PAC providers, approximately 16% of the $368 billion paid that year for Traditional Medicare (TM) claims (MedPAC, 2015a). This is similar to the $69 billion spending in the much-studied Medicare Part D program.

In recent years, the geographic variation and growth rate of spending on PAC have raised concerns about the efficiency of the sector. From 2001 to 2013, Medicare spending on PAC grew at an annual rate of 6.1%, 2 percentage points higher than the rate of total spending growth for TM (The Boards of Trustees for Medicare, 2002 and 2014; MedPAC, 2015a). A recent Institute of Medicine report found that, despite accounting for only 16% of spending, PAC contributed to a striking 73% of the unexplained geographic variation in spending, suggesting that there may be substantial inefficiencies in this setting (Newhouse et al., 2013).

It is useful to think about patients generally flowing "downstream" through the healthcare system. Upon an acute health event, they enter a regular, Acute Care Hospital (ACH), from there they may enter a PAC facility to recover, and eventually go home once they are sufficiently healthy and independent. Some ACH patients "skip" the PAC stay and return home directly from the ACH, and some patients occasionally move "upstream" from a PAC facility back to an ACH due to a relapse.

The top panel of Figure 2 gives a sense of transitions between ACHs, PAC facilities (LTCHs, SNFs, and IRFs), home (including home health agencies), and death (including hospice). In our data, described below, 26% of patients who are discharged from an ACH received follow-up care from a PAC facility (i.e. a SNF, IRF, or LTCH, but excluding HHAs).[4] From these PAC facilities, 60% of patients continue to flow home ("downstream"), where they may still receive treatment from an HHA, while 33% are discharged back upstream to an ACH. The remaining 7% of discharges are to a hospice or due to death.

Just like there is a natural flow of patients into and out of the PAC system, there is also a

[4] In analysis that includes HHAs in the calculation, the number of ACH patients who are discharged to PAC rises to 42% (MedPAC, 2015b).

general ordering of care within it. LTCHs provide the most intensive care, SNFs and IRFs provide less intensive care, and HHAs the least intensive bundle of medical services. For instance, the share of patients in the highest severity of illness (SOI) category declines from 43% at LTCHs, to approximately 12% at SNFs and IRFs, to 4% at HHA (AHA, 2010). Medicare payments per day follow the same declining pattern.

The bottom panel of Figure 2 looks at patient flows from LTCHs. About 11% of LTCH patients are discharged upstream to an ACH, 38% are discharged downstream to another PAC facility (SNF or IRF), and 33% are discharged to their homes, where they may continue to receive care from an HHA. The remaining 18% are discharged to a hospice (4%) or die within the LTCH (14%). In contrast, once in a SNF or IRF, patients almost never get discharged "upstream" to an LTCH, they die much less frequently (5%), and much more often (60%) return directly home.

Despite the interlocking nature of the PAC system, the way that Medicare reimburses facilities varies substantially by the type of facility. Historically, all facilities were paid according to an administrative estimate of their costs. Since the early 2000s, however, all types of PAC stays are paid under a prospective payment system (PPS), yet the unit of payment varies across sites. Loosely, HHAs are paid per 60-day episodes of care, SNFs are paid a fixed rate per day of stay, while IRFs and LTCHs are in principle paid a fixed amount per admission (like ACHs).[5] We provide more details on LTCH payments in Section 3.

The fact that each type of facility is paid under a different system has often raised concerns. From a public health perspective, there is concern that the separate payment systems do not give providers enough incentive to coordinated care across different facilities. From a budgetary perspective, there is concern that providers may shuffle patients across facilities with the aim of increasing Medicare payments. These concerns have spurred various proposals for payment reform, including a recent bill which proposes providing a "bundled payment" to a single PAC coordinator, and letting this coordinator internalize the costs and benefits associated with the sequence of admissions and discharges for the entire episode of care (H.R.1458 - BACPAC Act of 2015).

2.2 Long-Term Care Hospitals

Our primary focus is on patients whose point of entry into the PAC system is a long-term care hospital (LTCH). The demarcation "LTCH" describes how the provider gets paid by Medicare. For a hospital to get paid as an LTCH, it must have an average inpatient length of stay of 25 days or more. Naturally, there are many ways to meet this requirement, so from a medical standpoint the question of what exactly is an LTCH often results in a host of differentiated and fuzzy answers.

The LTCH category of hospitals was created to solve a potential problem created by the 1982 Tax Equity and Fiscal Responsibility Act (TEFRA), which established the prospective payment

[5] These different payment systems also has differential implications for beneficiaries' cost-sharing requirements across types of PAC facilities. Beneficiaries generally are not required for any cost sharing for HHA services. IRF and LTCH stays are tied to the beneficiaries' inpatient deductible, so when arriving from an ACH there would typically be no requirement for additional cost sharing. SNF stays are associated with a separate SNF deductible for stays longer than 20 days.

system (PPS) for acute care hospitals. Under the new prospective payment system, hospitals were paid per discharge and not based on their costs, as a way to create incentives for hospitals to be efficient in their treatment decisions. Regulators who were designing the PPS realized that there was a small number of hospitals that had long average length-of-stays (LOS) and would not be financially viable under the lump-sum PPS. LTCHs were thus created as a carve-out from acute-care hospital PPS for hospitals that had an average inpatient length of stay of 25 days or more; such hospitals were allowed to be paid according to a measure of costs, roughly in the spirit of the pre-1982 payment system. At that point in time, there were 40 hospitals in the U.S. that qualified as LTCHs. They were mainly former tuberculosis and chronic disease hospitals in the Boston, New York City, and Philadelphia metropolitan areas. Their payments were based on costs measured in 1982 and adjusted for inflation in subsequent years. See Liu et al. (2001) for more on the background of the LTCH sector.

Over the last 30 years, and perhaps because of the LTCH exception from PPS, there has been rapid growth of the LTCH sector. Because new entrants did not have cost data for 1982, payments for new entrants were determined by costs in the initial years of operation. This encouraged new entrants to be inefficient when they first opened and then earn profits by increasing their efficiency over time (Liu et al., 2001). From the initial 40 hospitals first designated as LTCHs in 1982, there are now over 400 such hospitals in the country.

Geographic penetration of LTCHs is extremely varied. This presumably reflects their historical roots as tuberculosis and chronic disease hospitals in the northeast, as well as certificate of need (CON) laws that have restricted entry. There are only a few LTCHs in the west of the country, and three states (Massachusetts, Texas, and Louisiana) account for a third of all LTCHs. Almost a quarter of Medicare PAC stays in the U.S. occur in hospital referral areas (HRRs) that have no LTCH. In places where there are LTCHs, these hospitals are an important part of Medicare's PAC landscape. For instance, in hospital service areas (HSAs) with at least one LTCH, we calculate that LTCHs account for 13% of Medicare PAC days and 28% of Medicare PAC spending.[6] Overall, payments to LTCHs account for 9.3% of Medicare PAC spending (MedPAC, 2015a)

LTCHs are much more likely to be for-profit than other medical providers. According to AHA data, 72% of LTCHs are for-profit (versus 17% for ACHs), 22% are non-profit, and 6% are government run. The LTCH market is dominated by two for-profit companies, Kindred Health Systems and Select Medical, which run about 40% of LTCHs. Company reports indicate that the LTCH business is highly profitable. For their business segments that include LTCHs, Kindred's profits have hovered between 22% and 29% of revenue and Select's profits have ranged between 16% to 22%.[7]

Approximately half of LTCHs are known as Hospitals-within-Hospitals (HwHs), meaning that they are physically located within the building or campus of an ACH but have a separate governing body and medical staff. Regardless of an LTCH's location (co-located or freestanding), they tend to

[6] The numbers we calculate exclude spending by HHAs.

[7] Profits are defined as EBITA. Kindred's profits are based on 2009 to 2015 company reports. Prior to 2009, Kindred did not separate out their reporting of LTCH profits from the much larger SNF category. Select's profits are based on company reports from 2004 to 2015.

have strong relationships with a single ACH.[8] Because of concerns over close relationships between LTCHs and their partner ACHs, in 2005 CMS established a policy known as the "25-percent rule" that creates disincentives for admitting more than 25% of patients from a single facility; however, Congress has delayed the full implementation of the law.[9]

2.3 Data

Our main data source is the Medicare Provider and Analysis Review (MedPAR) data, spanning the years 2000-2012. The dataset contains claim-level information on stays at ACHs, LTCHs, SNFs, and IRFs.[10] Each record is a unique stay for which a claim was submitted, and it contains information on procedures, admission and discharge dates, admission sources and discharge destinations, hospital charges, and Medicare payments. These data also provide us with basic demographic information such as the age, sex, and race of the beneficiary, and information about the patient's diagnoses (DRGs).

We supplement this primary source with several ancillary data sources. First, we use Medicare's beneficiary files to determine whether the beneficiary is dually eligible for Medicare and Medicaid and the date of death (if any). Second, we use the Medicare chronic conditions file to measure whether the individual has any of 27 chronic conditions. Third, we use data from the American Hospital Association (AHA) survey over the same period to determine whether a hospital is for-profit, non-profit, or government owned, and whether it is co-located with an ACH.

Our analysis focuses on the current Medicare payment schedule for LTCHs, known as LTCH-PPS. We analyze the time periods before and after full implementation of LTCH-PPS, which was phased in over a five year period starting on October 1, 2002. We define the *pre-PPS period* as January 1, 2000 to September 30, 2002. We exclude the October 2002 to September 2007 phase-in period because provider behavior during this period potentially reflects the combination of changing financial incentives and learning about the new incentive structure, complicating the interpretation of the data. We define the *PPS period* as October 2007 to September 2012.[11]

Table 1 shows summary statistics on ACH, LTCH, and SNF/IRF admissions in the pre-PPS and PPS periods.[12] Since an observation is an admission, some patients (16%) show up multiple

[8] MedPAC (2004) found that HwHs receive 61% of their cases from their most frequent referring hospital and freestanding hospitals receive 42% from their most frequent referring hospital.

[9] There is also a regulation known as the "5-percent rule" that addresses the incentive for HwH to "ping-pong" patients between the ACH and LTCH. In particular, if more than 5% of patients who are discharged from an LTCH to an ACH are readmitted to the LTCH, the LTCH will be compensated as if the patient had a single LTCH stay (42 CFR 412.532).

[10] Due to the scope of our data-use agreement, we do not have access to claim-level information about hospice stays or HHA services. Thus, for our main analysis of LTCH patients and their discharge decisions, we can observe all discharge destinations, but we cannot observe post-discharge claims for HHA or hospice.

[11] We begin our pre-PPS period in year 2000 because we do not have data from earlier years. We end the PPS period in year 2012 because we do not have more recent data.

[12] We group SNF and IRF admissions together for convenience, as both represent post-acute care that is "less intense" than an LTCH and because IRFs only account for a small (6.4%) fraction of these admissions.

times in the data. LTCH patients are, on average, slightly younger than ACH patients and much younger than SNF/IRF patients. LTCH patients are also almost twice as likely to be black and about one-third more likely to be eligible for Medicaid, relative to ACH and SNF/IRF patients. These differences are fairly stable over time. From a health perspective, the number of chronic conditions is slightly higher for both LTCH and SNF/IRF relative to ACH patients. Mortality rates are highest for LTCH patients, with about 15% dying within 30 days of admission and 30% dying within 90 days. These mortality rates are about 50% larger than mortality rates for SNF/IRF patients and twice as large as those for ACH patients.

In terms of the intensity of medical care, LTCH stays are closer to ACH stays than stays at a SNF/IRF. The majority of LTCH and ACH patients receive at least one medical procedure versus about 5% of patients who visit an SNF/IRF. The most common LTCH procedures (blood transfusion, and ventilation) is also more similar to those that occur at an ACH, relative to occupational and physical therapies, which are the most common procedures in SNF/IRF. Length of stay, however, is (by design) much more similar in LTCH to that of SNF/IRF. The average stay at an ACH is 5 days, while it is just over 25 days in LTCH and SNF/IRF.

The bottom rows of Table 1 show statistics on Medicare and out-of-pocket payments. Medicare payments in the PPS period average $2,087 per day at an ACH, $1,390 per day at an LTCH, and $507 per day at a SNF/IRF. However, because LTCH stays are much longer than ACH stays, per-admission Medicare payments at LTCHs average approximately $35,402, which is three times greater than per-admission ACH and SNF/IRF payments. Out-of-pocket payments at ACHs and LTCHs arise from Medicare's Part A deductible ($1,156 in 2012) and from co-insurance payments that apply when the patient has more than 60 hospital days in the benefit period ($289 per day in 2012). Because patients have no out-of-pocket exposure between the deductible and their 60th hospital day, out-of-pocket payments are a modest 7.7% of Medicare payments at ACHs and 5.5% at LTCHs in the PPS period. SNFs, on the other hand, have a separate co-insurance schedule with payments of $144.50 per day for stays in excess of 20 days, and have a much higher out-of-pocket share.

3 LTCH Payments, Discharge Patterns, and Outcomes

In this section we present descriptive analysis on LTCHs' response to financial incentives. We start by describing the LTCH budget set, including the large jump in payments that is our primary source of identification. We then show evidence on how discharge patterns and mortality rates vary with the budget set. This descriptive evidence motivates several of our key modeling choices in the dynamic model of LTCH discharge behavior, which we present in the next section.

3.1 LTCH Payments

We provide a basic overview here of how LTCH payments vary with the patient's length of stay, an object we refer to as the LTCH budget set or payment schedule. Appendix A provides much more detail. Figure 1 summarizes the payment schedules in the pre-PPS and PPS periods.

Prior to October 2002, LTCHs were paid their (estimated) daily cost, generating a linear relationship between the length of the hospital stay and payments. As described earlier, this "cost plus" reimbursement of LTCHs was seen as potentially encouraging inefficient entry into the LTCH market. Because of this and other concerns, the 1997 Balanced Budget Act (BBA) and 1999 Balanced Budget Refinement Act (BBRA) implemented a PPS for LTCHs. LTCH-PPS was phased in over a 5-year period starting on October 1, 2002 and was fully implemented by October 1, 2007. At a broad level, LTCH-PPS is designed to operate like the PPS for acute care hospitals (IP-PPS), under which hospitals are paid a lump-sum that is based on the patient's diagnosis (diagnosis-related group, or DRG) and does not vary with the patient's length of stay.

Much like LTCHs were originally created to address a potential problem with the introduction of PPS for ACHs, so do the details of the LTCH-PPS payment schedule can be thought of as attempting to address a potential problem arising from the introduction of PPS for LTCHs. In particular, in designing LTCH-PPS, officials were concerned that LTCHs might discharge patients after a small number of days but still receive large lump-sum payments intended for longer hospitals stays. To address this concern, they created short stay outlier (SSO) threshold. If stays were shorter than the SSO threshold, payments would be based on the pre-PPS cost-based reimbursement schedule and not a large lump sum. However, while reducing the incentive to cycle patients in and out of the LTCH, the SSO system creates potentially problematic incentives at the SSO threshold. At the day where payments switch from per-day reimbursement to lump-sum prospective payment amount, Medicare payments for keeping a patient an additional day "jump" by a large amount.

Figure 1 graphs the average payment schedules in the pre-PPS and PPS periods, pooling across LTCH facilities and DRGs. The y-axis shows cumulative Medicare payments, inflation adjusted to 2012 dollars. The x-axis shows the length of the stay relative to the SSO threshold, which we normalize to be day 0. The SSO threshold is defined as five-sixths the geometric mean length of stay for that DRG in the previous year and therefore varies by DRG and to a much lesser extent by year. The modal threshold (accounting for 22.4% of PPS stays) is 20 days and 99% of the sample has SSO thresholds between 16 and 39 days.[13] As a result, in this and subsequent figures, we present results relative to the SSO threshold so that we can pool analyses across DRGs.[14] Because the SSO threshold is undefined in the pre-PPS period, we assign pre-PPS stays the threshold for their DRG from PPS period.[15]

Under the pre-PPS system, average payments scale linearly with the length of stay at a rate of $1,071 per day. Under the PPS system, payments increase linearly by $1,386 per day to the left of

[13] Across all DRGs, the SSO threshold ranges from 14 to 56 days.

[14] We start the x-axis range at -15 days because nearly all SSO thresholds occur after 16 days. If we extended the x-axis range to -16, for example, there would be a change in the composition of DRGs between days -16 and -15 due to the entry of new DRGs into the sample. We end the x-axis range at 15 days because there are relatively few (14.6%) patients who are kept at the LTCH more than 15 days beyond the SSO threshold.

[15] In particular, since the thresholds occasionally vary over time, we use the SSO threshold from 2007, which is the first year in the PPS period.

the SSO threshold, "jump" by $13,656 at the SSO threshold, and remain constant thereafter. The increase in payments at the jump is large: it is equal to 55% of the cumulative payment amount on the day prior to the threshold, or equivalent to about 10 days of payments at the pre-threshold daily rate.

This sharp jump in payments was presumably not the intention of the policymakers who designed the LTC-PPS, but it arises naturally from the interaction of two sensible policies. As is standard in fixed price contracts, the LTCH-PPS payments were likely set to approximate average costs per stay. As noted, payments on a cost-plus basis up to the SSO threshold were introduced to avoid LTCHs receiving large lump sum payments for relatively short stays; under the pre-PPS payment scheme, a large share, 44%, of stays would have been below the subsequent SSO threshold, thus qualified as short stay "outliers." While reducing the incentive to discharge patients from an ACH to an LTCH, the SSO system unavoidably creates potentially problematic incentives at the transition from cost-plus to fixed-price payments for the stay. In Section 4 we explore the impact of alternative, counterfactual payment schedules, but short of getting rid of PPS, there is an unavoidable transition problem.

3.2 Discharge Patterns

We present a number of descriptive results on discharge patterns from the LTCH around the threshold. The results can be summarize as follows. First, there is a large spike in discharges at precisely the date of the jump in payments, indicating a strong response to financial incentives. Second, several pieces of evidence are consistent with the fact that the marginal patients discharged at the threshold are in relatively better health: they are disproportionately discharged "downstream" and they have lower mortality rates than patients discharged at other times. Third, among patients discharged downstream, the patient discharged at the threshold is relatively sicker, with higher post-discharge costs (paid to SNF/IRF or to ACH due to a relapse) than pre-threshold discharges.

Figure 3 shows the aggregate pattern of discharges by length of stay in the pre-PPS and PPS periods. A discharge occurs when the patient is transferred to another facility, sent home (possibly with home health care provided by an HHA), or dies at the LTCH. The y-axis shows discharges as a share of the total number of stays at the LTCH. The x-axis plots the length of stay relative to the DRG-specific SSO threshold, defined in the same manner as in Figure 1. In the PPS period, there is a sharp increase in discharges at the SSO threshold, with the share of discharges increasing from about 2% to 9% per day. Discharge rates remain elevated over the subsequent 7 to 10 days before reverting to baseline. In the pre-PPS period, there is no evidence of any bunching at the SSO threshold. This discharge pattern is consistent with a strong response to the financial incentives at the threshold.

Relative to the pre-PPS constant, per-day payment schedule, the increase in discharges under PPS at the threshold could be drawn either "from the left" or "from the right" of the distribution. In other words, the excess discharges at the threshold could reflect patients who under the pre-PPS schedule would have been discharged before the threshold but are "retained" in order to get the lump sum payout at the threshold, or patients who would have been discharged after the threshold but are now discharged earlier since there is no longer a marginal financial payment associated with keeping them additional days. Because the share of discharges to the left of the threshold is lower in the PPS period relative to the pre-PPS period, it is tempting to infer that the excess mass of discharges is primarily "drawn from the left" of the distribution. However, we caution that differences in the discharge rate might not only reflect the change in financial incentives but also changes in patient health and other secular trends between the pre-PPS and PPS periods. In fact, a simple reweighting of the pre-PPS admissions to match the DRG composition of the PPS period is sufficient to make the pre-PPS discharge rate much closer to that of the PPS discharge rate prior to the SSO threshold (Appendix Figure 1). In Section 4, we show how we can use our model to compare discharge patterns under the observed PPS payment schedule to discharge patterns under alternative, counterfactual payment schedules.

Figure 4 decomposes the discharge pattern by the location of discharge. For purposes of this exercise we group the discharges into three mutually exclusive and comprehensively exhaustive categories, in roughly descending order of patient health: discharges "downstream" (defined as SNF/IRF, or home with or without home health), discharges "upstream" (to ACH or hospice), and discharges due to death. The figure shows increases at the threshold in discharges both upstream and downstream, but the proportional increase is substantially larger on the downstream margin. Moreover, because the pre-threshold discharge rate is much higher downstream, the sharp change in discharge rate at the threshold (shown in Figure 3) is almost entirely driven by downstream discharges.

Appendix Figure 2 plots the 30-day mortality by length of stay, where the 30-day mortality rate is defined as death within 30 days of *discharge*. The graph shows a sharp drop in mortality at the SSO threshold, again suggesting that the patients who are discharged at the threshold are healthier than the patients who are discharged immediately beforehand. We caution that the decline in mortality not only reflects changes in the composition of patients discharged at the threshold, but could in principle reflect a treatment effect of discharge on health. We address this concern in the next section.

Figure 5 plots Medicare payments that occur after the LTCH discharge, by date of discharge. These post-discharge payments are defined as the sum of Medicare payments (to ACH, SNF, IRF, or LTCH) over the episode of care for patients discharged on different days. We define the episode of care as the spell of continuous days with a Medicare payment so that the episode ends if there are two days or more without any Medicare payments being made to any facility. We show separately the post-discharge costs for patients initially discharged "upstream" (to an ACH) and initially discharged "downstream" (to a SNF/IRF).[16] For patients initially discharged downstream to a SNF/IRF, there is a sharp increase in post-discharge payments at the threshold, with average post-discharge payments increasing from approximately $1,000 to $2,000. There is small change in the opposite direction in post-discharge payments at the threshold for patients initially discharged to

[16] In Figure 4, "upstream" discharges include discharges to hospice and "downstream" discharges include discharges to home health, but we do not observe payments to either of these in our data and so exclude them from this analysis.

an ACH.

Figure 5 suggests a simple model of the behavioral response to the threshold, which motivates the model we present in Section 4. Prior to the threshold, retaining patients is profitable, and only the healthiest patients are discharged to SNF/IRF and the sickest patients discharged to an ACH. After the threshold, on the SNF/IRF margin, LTCHs work "down the distribution" and discharges less healthy patients, increasing the post-discharge costs on average. Similarly, on the ACH margin, LTCHs work "up the distribution," discharging patients who are in better health, and decreasing average post-discharge costs. The marginal patient discharged downstream at the threshold is therefore sicker than the average patient discharged prior to the threshold, while the marginal patient discharged upstream is slightly healthier than the average patient discharged upstream in prior days.

3.3 (Lack of) Mortality Effects

A natural question raised by the discharge patterns is whether the distortions in the timing of discharges have an impact on patient health and in particular mortality. Since the 90-day mortality rate of LTCH patients is approximately 30%, if these distortions are harmful to health, it seems plausible that we might be able to pick up an effect with our data. To guide the interpretation of the mortality results, it is helpful to think about health evolving according to a stochastic process, with sicker patients having a higher probability of death. Difference in the location of care might impact the level of someone's health, generating an on-impact effect on the probability of death. Changing locations might also effect the stochastic process for health, which would be associated with a longer-run change in mortality rate, but might not have an immediate, on-impact effect.

We have already seen in the bottom panel of Figure 4 some suggestive evidence that mortality rates are declining over the course of the LTCH stay (which is not surprising given natural selection; as the sickest die, the remaining patients are gradually healthier) with little difference around the SSO threshold. However, the interpretation of the bottom panel of Figure 4 – which plots mortality rates for LTCH patients by length of stay – is complicated by selection concerns. Since LTCHs are differentially discharging healthier patients at the SSO threshold, the composition of patients who remain at the LTCH is changing, making it tricky to disentangle any potential treatment effects on mortality from effects due to changes in the selection of LTCH patients.

In Figure 6 we circumvent this concern by taking advantage of the fact that our data allow us to track mortality outcomes for patients even after their LTCH discharge. Conceptually, our mortality analysis follows the logic of a reduced form regression, where the mortality hazard is the outcome, discharge patterns are the endogenous variable, and the threshold in the PPS payment schedule is the instrument. In particular, since we know there is a sharp jump in discharge patterns at the threshold (analogous to a large first stage), if there is also a change in the level or slope of the mortality hazard in the vicinity of the threshold (that is, non-zero reduced form), we can infer that the distortion in discharge location has an impact on mortality.

The top panel of Figure 6 is thus similar to the bottom panel of Figure 4, but uses the full set of LTCH patients (unconditional on their location) rather than only those who have yet to be discharged. As before, natural selection leads mortality rates to decline over time, but we now can interpret more cleanly the mortality pattern around the SSO threshold. The plot shows no obvious evidence of a change in the level of mortality hazard in the vicinity of the threshold during the PPS period, with perhaps a small shift of mortality from just before the threshold to just after. The evidence is consistent with no mortality effect but it does not allow us to rule out a gradual effect that would not appear sharply in the data.

The bottom panel of Figure 6 attempts to look for a more gradual effect by plotting a 30-day mortality rate by days since LTCH admission. If distortions in the location of care effected the stochastic process for health, we might not observe an immediate effect, but would see a change in mortality over a longer time horizon. The 30-day mortality hazard measures the share of patients who are alive on a given day but die in the next 30 days, allowing us to observe a potential longer-run effect. The plot shows no effect around the threshold, suggesting there are no gradual effects of the distortions in discharges on mortality. Obviously, this (lack of) reduced form effect needs to be judged in relation to the size of the first stage effect on the location of care. And as we show in our counterfactuals, the "experiment" we analyze only shifts the location of care for a relatively small number of days, so perhaps the non-effect is not surprising. Yet, these relatively small changes in the location of care are precisely what we explore with our counterfactuals. Figure 6 makes us conclude that there is little evidence of quantitatively large effect on mortality that is created by the sharp discharge incentives at the SSO threshold.

We can gain additional traction on potential mortality impact by exploiting difference in the response to incentives between for-profit and non-profit LTCHs. Figure 7 shows the payment schedules (top panel), discharge shares (middle panel), and mortality rates (bottom panel), separately by for-profit status. While the payment schedules are almost the same across the two groups of hospitals, the behavioral response to the jump in payments is substantially larger for for-profit LTCHs. Prior to the SSO threshold, for-profit hospitals have a slightly lower discharge rate, but at the SSO threshold their discharge share rises by about twice as much as that of non-profit hospitals, suggesting that, perhaps not surprisingly, for-profit hospitals are relatively more responsive to financial incentives.

Given the for-profits LTCHs' much larger behavioral response to the jump in payments, if the PPS payment schedule affected mortality, we would expect a more pronounced effect at for-profit hospitals. The bottom panel of Figure 7 provides no suggestion of any differences in the mortality pattern by for-profit status. Mortality hazards are remarkably similar across both groups, both in the immediate vicinity of the SSO threshold – suggesting no on-impact effect of the location distortion on mortality – and also in the weeks before and after the threshold – suggesting that the distortion in discharge behavior does not have a less-immediate, gradual effect on mortality.[17]

Overall, we interpret the results in this section as showing no evidence of any impact of the

[17] The mortality pattern for non-profit hospitals is slightly more noisy, but this is presumably driven by sampling variation due to the smaller samples size of non-profit admissions (113,154 versus 487,988 for for-profits in the PPS period).

PPS payment schedule relative to the pre-PPS schedule on mortality. While these results are not definitive – given the challenges discussed in detecting delayed mortality effects – they provide no "smoking gun" evidence of patient harm (at least as measured by mortality). Combined with the earlier results, which indicate that the patients who are most affected by the SSO threshold are disproportionately healthy, the results suggest that the marginal patient affected by the PPS payment schedule is able to receive similar care whether they are located in an LTCH or in their alternative setting, which empirically is usually a SNF.

4 Quantifying the Importance of Financial Incentives

The results in the last section provide descriptive evidence of the response of LTCHs to the sharp financial incentives associated with the SSO threshold. One way to quantify the importance of the financial incentives in directing discharge patterns out of LTCHs is to assess how these patterns would change in response to counterfactual financial contracts that exhibit weaker incentives. Doing so requires a dynamic model, which is the focus of this section.

4.1 Model and Estimation

4.1.1 A simple model of LTCH discharge decision

Consider a patient i who is admitted at day $t = 0$ to LTCH l. We index patient i's health at the time of admission by $h_{i,0}$, and assume that $h_{i,t}$ (conditional on patient i staying at LTCH l) evolves stochastically from day to day. In particular, we assume that $h_{i,t}$ follows a monotone Markov process, such that $h_{i,t} \sim F(\cdot|h_{i,t-1})$ with $F(\cdot|h)$ stochastically increasing in h. We use higher values of h to indicate better health and thus assume that daily mortality hazard $m(h)$ is strictly decreasing in h.

Hospital l's flow (daily) payoff from keeping patient i (whose health is given by h) during the tth day since admission is given by

$$u_l(h, t|keep) = p(t) - c_l(h) + \beta w_l(h) + \epsilon_{ilt}, \tag{1}$$

where $p(t)$ is the hospital's revenue, which depends on CMS' reimbursement schedule for patient i, $c_l(h)$ is the hospital's daily cost of treating a patient with health index h, and the third term captures the patient's utility from staying at LTCH l, $w_l(h)$, multiplied by the hospital's weight on it β. Finally, ϵ_{ilt} is an error term, distributed i.i.d. type I extreme value, presumably capturing idiosyncratic considerations associated with the patient and/or the hospital.

Our focus is on the hospital's discharge decision. Consider a set of J alternative destinations for patient i, each indexed by j. Conditional on discharging the patient to destination j, LTCH l's revenue and cost are both zero, and its flow payoffs are given by the patient's utility, again multiplied by the hospital's weight on it β:

$$u_j(h, t|discharge\ to\ j) = \beta w_j(h) + \epsilon_{ijt}. \tag{2}$$

Moreover, because hospital l loses control over the patient upon discharge, it will be convenient to denote by $W_j(h)$ the present value of the patient's utility from being discharged to alternative j.

This setting lends itself to a simple dynamic programming problem, which can be represented by the following Bellman equation:

$$W_l(h,t) = E\left(max \left\{ \begin{array}{c} u_l(h,t|keep) + \delta\left(1-m(h)\right)\int W_l(h',t+1)dF(h'|h), \\ max_{j\in J}\left(u_j(h,t|discharge\ to\ j) + \delta\left(1-m(h)\right)\int W_j(h')dF(h'|h)\right) \end{array} \right\} \right), \tag{3}$$

where δ is the LTCH's (daily) discount factor. That is, the two state variables are the health of the patient and the number of days since LTCH admission. Every day the hospital makes a binary decision whether to discharge or keep the patient, and in the event of a discharge the hospital also decides about the discharge location. Of course the model can allow the patient to actually "decide" about the discharge destination by having the hospital place a large weight on patient utility.

It is convenient to benchmark the patient's utility against her utility at the LTCH, thus normalizing $w_l(h) = 0$ for all h, $v_j(h) = w_j(h) - w_l(h)$, and $V_j(h)$ is defined accordingly. Applying these adjustments and using the well-known expression for the logit's inclusive value, we can write the problem as

$$V_l(h,t) = ln\left\{ exp\left(p(t) - c_l(h) + \delta\left(1-m(h)\right)\int V_l(h',t+1)dF(h'|h)\right) + \sum_{j\in J} exp\left(V_j(h)\right)\right\}. \tag{4}$$

Finally, we note that the state variable t only affects the problem through the hospital revenue function $p(t)$, and $p(t) = 0$ for all $t > SSO$, so the problem becomes stationary after the SSO threshold, and the solution is simply a fixed point of

$$V_l^{t>SSO}(h) = ln\left\{ exp\left(-c_l(h) + \delta\left(1-m(h)\right)\int V_l^{t>SSO}(h')dF(h'|h)\right) + \sum_{j\in J} exp\left(V_j(h)\right)\right\}. \tag{5}$$

We can therefore solve for the dynamic problem by first solving for the fixed point associated with the post-SSO stationary part of the problem given by equation (5), and then iterating backwards until $t = 0$ using equation (4).

4.1.2 Parameterization, estimation, and identification

Parameterization. We make several additional assumptions before we take the model to the data. First, we restrict the set of alternative discharge destinations J to include only two options, $J = \{S, A\}$. Motivated by the summary statistics described in Section 3, option S covers a collection of downstream destinations – SNF, IRF, home care, and "other" – that are appropriate for LTCH patients who are of better health or require lower levels of medical monitoring. In contrast, option A covers upstream discharge destinations – ACH and hospice – which would be natural discharge destinations for patients who are of worse health.

Because, conditional on discharge, financial incentives do not affect LTCH's discharge destination, having a richer set of discharge options is unlikely to affect our counterfactual. By focusing on two options, we essentially restrict the LTCH to consider two types of marginal LTCH patients. One set of marginal patient are those who are healthier, and for whom the hospital must consider whether to keep them or discharge them to location S. The second set of marginal patients are sicker, and for whom the hospital must consider whether to keep them in the LTCH or transfer them to location A.

The second assumption regards the health process. Given that mortality is monotone in h, it is convenient to normalize the health index by mortality risk. We do so by assuming that h is defined by its associated mortality hazard using the following relationship:

$$m(h) = 1 - \Phi(h), \tag{6}$$

where $\Phi(\cdot)$ is the standard normal CDF. We note that h is an index and thus has no cardinal meaning, and the above is simply a normalization, which entails h with a cardinal measure. Equipped with this normalization, we then make parametric assumptions about the initial (as of LTCH admission, $t = 0$) health distribution of newly admitted patients, and about how the health process evolves over time. Specifically, we assume that $h_{i,0}$ is drawn from $N(\mu_0, \sigma_0^2)$ and that $F(\cdot|h_{i,t-1})$ follows a simple AR(1) process:

$$h_{i,t} = \mu + \rho h_{i,t-1} + \varepsilon_{i,t}, \qquad \text{where} \qquad \varepsilon_{i,t} \sim N(0, \sigma^2). \tag{7}$$

In our baseline specification, we allow the health process to be different in the pre-PPS and PPS periods to accommodate potential differences in the LTCH patient mix, which may result from the growth of the LTCH sector, time trends in medical technology and practice, or directly from the change in financial incentives.

The third assumption is associated with the LTCH's cost, $c_l(h)$, which we assume are given by

$$c_l(h) = \gamma c^{reported}. \tag{8}$$

That is, as described in Section 2 and reported in the bottom row of Table 1, we observe the cost associated with each hospital l, which enters the formula by which it is paid by CMS. We do not treat these reported costs as the "true" costs, but we use it to guide our model of costs in two ways. First, and importantly, we assume that the hospital's cost do not vary with the patient health h, which is consistent with CMS's treatment of costs and also seems natural given that LTCH patients are generally stable. Second, we assume that the reported costs are true up to a monotone transformation, which we assume to be linear. This assumption means that if hospital l has higher reported costs than hospital k, we will also assume that this also reflects the ranking of their true underlying cost. This seems natural, and could be driven by a variety of factors, including geographic location. We would naturally expect $\gamma \leq 1$.

The fourth assumption is to parameterize $V_A(h)$ and $V_S(h)$. We approximate each using a linear function in h, so that

$$V_j(h) = v_{0j} + v_{1j} h \qquad \text{for} \qquad j = A, S. \tag{9}$$

Recall from Section 3 that healthier patients (higher h), who are associated with lower mortality, are discharged to S, while sicker patients (lower h), associated with higher mortality, are discharged to A. It is therefore natural to expect $v_{1,S} > 0$ and $v_{1,A} < 0$. That is, all else equal, facility S becomes a more attractive discharge destination as health gets better (h is higher) and facility A becomes a more attractive discharge destination as patients' health worsens (h is lower). As explained below, one of the intercept terms $v_{0,S}$ and $v_{0,A}$ needs to be normalized, so we set $v_{0,A} = 0$.

Finally, as is typical in these type of models, we set (rather than estimate) the daily discount factor to $\delta = 0.95^{1/365}$. Thus, overall we are left with 14 parameters to estimate: five parameters $(\mu_0, \sigma_0, \mu, \sigma, \rho)$ that are associated with the health distribution and the way it evolves over time in the pre-PPS period, five corresponding parameters in the PPS period, the cost parameter (γ), and the three parameters (v_{1A}, v_{0S}, v_{1S}) associated with the relative value of patients at facilities A and S.

Estimation. An important decision is how to treat heterogeneity across patients, observable health conditions, and LTCH hospitals. In our baseline specification, we abstract from such heterogeneity and instead model the "average" discharge decision as it pertains to the "average" LTCH patient and the "average" payment schedule. That is, we pool all payment schedules observed in the data, separately for the pre-PPS and PPS periods, measure each day in the schedule relative to the DRG-specific SSO threshold in the PPS period (which is normalized to zero), and construct the average payment schedule for each day, as shown in Figure 1. We then apply an analogous exercise to the discharge pattern, and construct the distribution of discharge patterns in a 31-day window around the SSO threshold, as shown in Figure 3 and Figure 4. We then estimate our model in an attempt to match these average patterns. The dramatic difference in the payment schedules between the pre-PPS and PPS periods will assist in the identification of some of the parameters and is an important ingredient in our research design. An advantage of this approach of focusing on the average pattern rather than the heterogenous pattern is that it only requires us to solve the dynamic problem once (for each pricing period), which is computationally attractive. In future work, we plan to explore specifications that are less restrictive (e.g., by partitioning the data into bins based on the number of days to the SSO threshold and by conditioning on early-day mortality rates).

We estimate the model using simulated method of moments, by trying to match the daily mortality and discharge patterns presented in Figure 4. Specifically, we use 93 moments for the pre-PPS payment schedule, reflecting the daily discharge and mortality risks within a 31-day window around the SSO threshold. One set of moments is associated with discharge rates to S, another with discharge rates to A, and a third with mortality rates. We then construct another set of 93 corresponding moments for the PPS period. Because much of the identification is driven by the sharp change to discharge incentives at the SSO threshold, we assign greater weights to moments that are closer to day zero (the SSO threshold) by making weights linearly decline for days that are further away from day zero, hitting one-half of the day-zero weight at 15 days before or after the SSO threshold.

Generating the model predictions requires us to solve the dynamic problem described in the previous section for each set of parameter. To ease with computation, we approximate the health process $F(\cdot|h_{i,t-1})$ using a discrete space of health rates, that evolve according to a Markov transition matrix (Tauchen, 1986). This eases the solution of the dynamic problem, and at the same time allows us to read the discharge probabilities directly off the solution, without any need to integrate (presumably by simulation) over unattractive integration regions. The Appendix provides more details.

Intuition for Identification. To see the intuition for the identification of the parameters, it is easiest to consider first a case where health is homogeneous across patients and over time. Under this assumption, the data can be characterized by daily observations of discharge shares to A and to S (s_{At} and s_{St}, respectively), with the remaining patients staying at the LTCH ($s_{Lt} = 1 - s_{At} - s_{St}$). The problem resembles a repeated static discrete choice problem, where the mean utility of each discharge destination is given by the continuation values V_{At}, V_{St}, and V_{Lt}. As is usual in multinomial logit models, the observed daily shares can then be inverted to recover the values of V_{At}, V_{St}, and V_{Lt}, subject to a required level and scale normalizations.

Let us start with the level normalization. Setting $V_{At} = 0$ allows us to estimate V_S from the average values of V_{St}, and $V_{Lt,}$ up to a scale normalization. To identify the scale parameter, recall that V_{Lt} varies over time due to the expected present value of payments and cost. In particular, the present value is function of the payment schedule $p(t)$; reported costs scaled by parameter γ; and the relevant time horizon, which depends on the mortality rate and subsequent, endogenous discharge decisions. So one can think of the identification of the scale parameter and γ as a projection of the values of V_{Lt} on these observables. The sharp change in payments at the SSO threshold provides a sharp change in the present value of payments and identifies the scale parameter (or equivalently the coefficient on payments when the variance of the error term is standardized), and the differential change in the present value of payments versus costs as the patient approaches the SSO threshold identifies γ. This identification can be achieved from the PPS moments alone, but given that we restrict these parameters to be time-invariant, it is also aided by variation in discharge patterns between the pre-PPS and PPS periods.

If health status h was observed, we could make the argument above conditional on health, and thus identify each object as a non-parametric function of h. In practice h is unobserved, but identifying the health process is conceptually easy given our assumptions. If there are no discharges, which is roughly the case during the first week or so of the LTCH stay, the only attrition from the sample is due to mortality. With only five parameters that determine the initial health distribution and how it evolves from day to day, mortality rates over five days are sufficient to identify the health process parameters, separately in the pre-PPS and PPS periods. Once the unobserved health distribution is identified, we can integrate over h and apply a similar intuition to the one we described above for the homogenous h case. Moreover, once the health process is identified, the cross-sectional distribution of h varies over time in "known" ways, so we can also identify how the key parameters – in particular the V's – vary as a function of h.

Obviously, as is typically the case, the intuition for identification requires us to have substantial variation in the data. In practice, some of the variation is not as large, and statistical power issues require us to impose more parametric structure, so the estimable model is not as flexible – especially in terms of the extent to which parameters vary with h – as the identifiable structure would be.

4.2 Results

4.2.1 Parameter estimates and model fit

Table 2 presents the parameter estimates. We estimate $\gamma = 0.93$ implying that LTCH's actual costs are 7% lower than their reported value. This is consistent with our prior that reported costs are somewhat inflated.

The $v_{1,A}$, $v_{0,S}$, and $v_{1,S}$ parameters capture the value the LTCH places on the patient's utility from being discharged to A or S relative to remaining at the LTCH. The estimates imply that LTCHs are indifferent between A and S for a patient with $h = 1.75$, which is a fairly low health level. For instance, $h = 1.75$ is the 9.8th percentile of the steady state PPS health distribution ($\mu = 4.27$, $\sigma = 1.95$) and corresponds to a daily mortality hazard of 3.95 percent. Consistent with our description of patients flowing "downstream" as their health improves, S is relatively better for healthier patients and A is better for sicker patients. The magnitude of the slope parameter $v_{1,S}$ is about one-sixth as large (in absolute value) as the slope parameter $v_{1,A}$, which indicates that a given change in financial incentives will have a much larger effect on discharges on the downstream S margin. These estimates are consistent with the descriptive evidence that shows a substantially larger response on the downstream margin at the SSO threshold.

We are cautious not to over-interpret the health process parameters. Because they are the only parameters that are allowed to vary across the time periods, they capture not only differences in the health of admitted patients but may also reflect other factors that vary over time, such as changes in medical technology or the administrative capacity of providers.

The model fits the data reasonable well. Figure 8 presents our moments and the simulated moments from the estimated model. The left column shows values in the PPS period and the right column shows values in the pre-PPS period. The top row shows the share of discharges to A by day relative to the SSO threshold, the middle row shows the share of discharges to S, and the bottom row shows the share of patients who die at the LTCH. The model does a very good job fitting the "spike" in discharges to A and S in the PPS period. This is particularly encouraging because this variation is our key source of identifying variation. The model fit for the mortality patterns in the pre-PPS and PPS periods is good over the initial days, but less good at longer time horizons. This is likely due to our fairly parsimonious parameterization of the health process. The model fit is also poorer for discharges to A in the pre-PPS period.

Figure 9 provides some intuition for how the model operates. The solid black lines show the LTCH's discharge policy function at the estimated parameters.[18] Healthy patients (lower h) are

[18] The discharge policy function is not a deterministic function of h; given the ϵ's in the model, h is related to discharge stochastically. The policy lines in Figure 9 are drawn so that at that given level of h, 50% of the patients

discharged to S, while sick patients (higher h) are discharged to A. Consistent with the descriptive evidence, LTCHs work "down the distribution" at the jump and lower their discharge threshold on the S margin and conversely work "up the distribution" on the A margin and increase the discharge threshold. The larger shift on the S margin relative to the A margin relates directly to our discussion above on the magnitude of the slope parameter estimates ($v_{1,S}$ and $v_{1,A}$). The relatively small outward shift in the policy function just before the SSO threshold is consistent with the descriptive results which show limited evidence on "missing mass" immediately to the left of the SSO threshold.

The dashed lines in Figure 9 correspond to the different health trajectories for a new LTCH patient whose health status at admission is at the median of the distribution. For such a patient, we simulate health status forward (including death, which is an absorbing state), and plot the percentiles of this distribution. The cross sectional health distribution is reasonably stable at better health, but due to mortality is gradually deteriorating if the patient is sicker. The flat pattern of the percentiles for healthier patients helps explain what happens at the SSO threshold: the health distribution is not very different before and after the threshold, but the sharp change in the policy function means LTCH discharge patients, which otherwise would have been kept at the LTCH, to SNFs and other "downstream" locations.

4.2.2 Spending under counterfactual financial incentives

We use the model to examine how patient flows and total spending are affected by alternative payment schedules. Trying to stay close to the observed variation, we consider payment schedules which are "capped" in the sense that marginal payments are (eventually) reduced to zero. These schedules are in the spirit of the current PPS system, which caps payments at the PPS amount but do not generate the perverse financial incentives caused by the jump in payments. Under the capped payment schedules we consider, most patients are discharged within 30 days of admission, meaning that our estimates are not excessively sensitive to assumptions on the health process at longer time horizons.

Figure 10 compares the observed payment schedule to the three counterfactual payment schedules we consider. The top panel shows a counterfactual schedule we call "no jump, lower cap," which simply eliminates the jump in payments at the SSO threshold by reducing the PPS payment. This schedule is less generous than the current schedule: holding discharge patterns unchanged, it pays LTCHs the same amount for patients discharged before the SSO, but pays them less for patients who are discharged at the SSO or after. The middle panel shows a counterfactual we call "no jump, extended SSO threshold." This schedule is still less generous than the baseline schedule but more generous than the first counterfactual: it continues to pay LTCHs the per-day amount after the SSO until the total payment hits the PPS value. The bottom panel shows a counterfactual we call "no jump, higher rate per day," which eliminates the jump by paying the LTCHs at a higher daily rate leading up to the SSO threshold. This schedule is more generous than the current schedule:

are discharged to the relevant destination.

holding discharge patterns unchanged, it pays LTCHs the same amount for patients discharged at the SSO or after, but pays them more for any patient who is discharged prior to the SSO threshold.

We use our model to simulate discharge decisions and Medicare payments under each of these three counterfactuals. In particular, we assume that the initial distribution of health of admitted patients stays the same but that the subsequent discharge decisions reflect the incentives provided by the counterfactual payment schedules. Figure 11 shows the discharge patterns under each schedule and Table 3 summarizes the impact of these payments schedules on Medicare payments to the LTCH and to other facilities.

This latter aspect – payments to other facilities – is not part of our model, and requires some elaboration prior to our discussion of the results. For our counterfactual analysis, we are interested in the effect of alternative payment schedules on government costs for the entire episode of care, which includes the LTCH stay but also post-LTCH stays at other facilities within and outside the PAC system. For instance, if a less generous payment schedule induces LTCHs to more quickly discharge patients to an SNF, it would decrease LTCH payments. However, to the extent that these patients are sicker than the typical SNF patient, they may have longer SNF stays, generating an offsetting increase in Medicare payments to SNFs. To account for these potentially offsetting effects, we construct data on the average post-discharge cost P_{jt} by discharge day t and destination $j \in A, S$,[19] and project $\ln P_{jt}$ on $\overline{h}_{jt}$, where $\overline{h}_{jt}$ is the average health status (as predicted by the estimated model) of patients discharged to destination j on day t. We find, naturally, that sicker dischargees are associated with somewhat higher post-discharge costs, thus capturing potential offset effects. See the Appendix for more details.

As shown in Figure 11, under the "lower cap" payment schedule, the elimination of the jump in payments induces LTCHs to discharge a much higher share of patients before the SSO threshold, with the daily share of discharges to S increasing four-fold and discharges to A increasing more modestly over most of the pre-threshold period. The pool of remaining patients becomes healthier on average as indicated by the lower share of patients who die at the LTCH. Table 3 (column (2)) shows that under this counterfactual, the average length of stay at the LTCH is reduced from 18.2 to 13.5 days,[20] and Medicare payments to the LTCH are reduced by \$11,751 or 45 percent. We decompose this reduction in payments into a "mechanical" effect, calculated as Medicare payments under the counterfactual payment schedule holding discharge patterns constant at their baseline levels, and a behavioral response. The mechanical effect of the "lower cap" payment schedule is a reduction in payments of \$8,744 or about 75% of the overall reduction, with the remaining 25% due to the behavioral response of LTCHs to the counterfactual incentives.

The remaining rows of Table 3 consider the impact of this counterfactual payment schedule on Medicare payments throughout the rest of the episode of care. The counterfactual payment schedules can influence post-LTCH Medicare payments for two reasons. First, under the counter-

[19] This is the same as the data used to generate Figure 5.

[20] Length of stay is measured from day -15. To make it comparable to the summary statistics reported in Table 1, both numbers should be increased by 7.5 days (because the average SSO threshold across admissions in our sample is 22.5 days).

factuals, patients may be discharged to different locations with different implications for Medicare spending. For example, a patient whose health is improving and would be retained until the SSO threshold and discharged to a SNF under the baseline schedule might be immediately discharged to an ACH under a less generous counterfactual. Second, holding discharge location fixed, patients may be discharged on different days and at different health levels, also affecting Medicare spending. For instance, a patient whose health is declining might be discharged to an SNF at the SSO threshold under the baseline schedule but discharged to an SNF at an earlier date under a less generous counterfactual. Because SNF payments depend on health at admission, Medicare might pay less under the counterfactual because the admitted patient was in better health.

Under the "lower cap" counterfactual, there is a small increase in the share of patients discharged to A and a small change in payments per patient discharged to this location, consistent with LTCHs "moving up" the health distribution in their discharge decisions. On net, these effects raise post-LTCH costs by a modest \$750. On the S margin, there is a 4 percentage point increase in the share of discharges but also a \$1,000 decrease in payments per discharge. The decline in payments results from LTCHs discharging patients with declining health at shorter lengths of stay, when they are in better health, and their associated SNF payments are lower. The effects almost perfectly offset each other, yielding virtually no effect to post-LTCH cost.

The "extended SSO" counterfactual is more generous than the first counterfactual but still less generous than the baseline schedule. Under this schedule, discharge patterns to S and A fall roughly between those under the first counterfactual and the baseline. Table 3 shows that average length of stay is actually longer than the baseline, as LTCHs have an incentive to retain patients to day 9 (rather than day 0), when the per-day payments are capped. However, Medicare payments decrease by \$4,348, which is about half the size of the decrease under the first counterfactual. The offsetting nature of this counterfactual payment schedule makes the average behavioral response more similar to the observed one, so most of the decrease (88%) in LTCH payments is driven by the mechanical effect while only 12% is due to the behavioral response. Post-LTCH payments for patients discharged to A are very similar to those under the baseline schedule. Post-LTCH costs for patients discharged to S decline by almost \$2,000, primarily due to reduced payments per discharge. Thus for this counterfactual, accounting for post-LTCH payments is important, raising Medicare savings to \$5,926 or about 36% more than the direct savings from reduced payments to LTCHs.

The "higher rate per day" counterfactual, while eliminating the jump at the SSO threshold, is more generous than the baseline payment schedule. Under this schedule, discharge rates to S and A are similar to the baseline discharge patterns in days -15 to -7. Because of the elimination of the jump, however, there is less incentive to retain patients in the days immediately before the SSO threshold. Because of the reduced incentive to hang onto patients, average length of stay under this schedule is about 2 days shorter than under the baseline. LTCH payments, perhaps surprisingly, are only \$560 or 2.1% higher than under the baseline schedule. This is due to offsetting mechanical and behavioral effects. Holding discharge patterns fixed, applying this schedule would increase Medicare payments by \$1,588, which is almost 3 times larger then the overall increase, implying that the behavioral effect is a reduction in Medicare payments of \$1,028. Post-LTCH Medicare payments are almost identical to the baseline values.

4.2.3 For-profit vs. not-for-profit LTCHs

In the end of Section 3, we presented descriptive evidence where we split the sample by whether the LTCH was a for-profit or non-profit hospital. In particular, we showed that while payment schedules were similar for for-profit and non-profit LTCHs, for-profits exhibited a larger increase in discharges at the SSO threshold, suggesting that these facilities are more sensitive to financial incentives. However, we cautioned against drawing too strong conclusions from the descriptive results, as this finding might, in part, reflect differences in the health of patients admitted to these facilities. Our model allows us to adjust for differences in health, allowing us to more precisely isolate heterogeneity in the response to incentives.

To conduct this analysis, we estimate the model separately on the set of for-profit and non-profit LTCHs. Appendix Table 1 shows parameter estimates from these separate runs of the model. Our estimate of costs is slightly lower for for-profit ($\gamma = 0.88$) versus non-profit ($\gamma = 0.92$) hospitals, which is consistent with slightly higher operational efficiency by the private sector. For for-profit hospitals, the slope parameters on patient preferences ($\upsilon_{1,H}$ and $\upsilon_{1,S}$) are approximately one-quarter smaller in magnitude, indicating that for-profits are relatively more elastic in their discharge decisions to changes in financial incentives.

Comparing for-profits and non-profits parameter by parameter is unsatisfying. To quantify the difference between for- and non-profits over the full set of parameters, Table 4 conducts the counterfactual of assigning every LTCHs for-profit versus non-profit status. For reference, column (1) shows outcomes under the baseline organizational form. To construct column (2), we assign all LTCHs the γ, $\upsilon_{1,A}$, $\upsilon_{0,S}$, and $\upsilon_{1,S}$ estimates for for-profit hospitals but maintain the health process parameters that differ by organization formal so that there is "no change" in the health distribution of the patients. For column (3), we do the same but assign all LTCHs the non-profit estimates for γ, $\upsilon_{1,A}$, $\upsilon_{0,S}$, and $\upsilon_{1,S}$ holding the health process fixed. Appendix Figure 3 shows discharge patterns under these counterfactuals.

Consistent with the increased responsiveness to incentives, Medicare payments are \$1,728 (or about 7%) higher under the for-profit counterfactual than under the counterfactual where we assign every hospital non-profit status. Comparing the non-profit counterfactual to the baseline, the estimates imply that "converting" the current set of for-profits to non-profit status would reduce payments by \$1,351 or about 5% of their current level, holding the distribution of patient health fixed.

4.2.4 Accounting for the endogeneity of PAC patient mix

TBA

5 Conclusions

TBA

References

American Hospital Association (AHA). 2010. "TrendWatch: Maximizing the Value of Post-Acute Care." Available at http://www.aha.org/research/reports/tw/10nov-tw-postacute.pdf.

Boards of Trustees for Medicare. 2002. "The 2002 Annual Report of the Boards of Trustees of the Federal Hospital Insurance Trust Fund and the Federal Supplementary Medical Insurance Trust Fund."

Boards of Trustees for Medicare. 2014. "The 2014 Annual Report of the Boards of Trustees of the Federal Hospital Insurance Trust Fund and the Federal Supplementary Medical Insurance Trust Fund."

Bureau of Economic Analysis (BEA). 2016. "Download NIPA Tables: Government Current Receipts and Expenditures (Section 3)." National Economic Accounts. Available at http://www.bea.gov/national/nipaweb/DownSS2.asp. Accessed August 31, 2016.

Chetty, Raj, John Friedman, Tore Olsen, and Luigi Pistaferri. 2011. "Adjustment Costs, Firm Responses, and Micro vs. Macro Labor Supply Elasticities: Evidence from Danish Tax Records". *Quarterly Journal of Economics* 126 (2): 749-804.

Clemens, Jeffrey and Joshua D. Gottlieb. 2014. "Do Physicians' Financial Incentives Affect Medical Treatment and Patient Health?" *American Economic Review* 104(4): 1320-1349.

Congressional Budget Office (CBO). 2016. "Updated Budget Projections: 2016 to 2026," Congressional Budget Office Publication #51384.

Cutler, David, and Richard Zeckhauser. 2000. "The Anatomy of Health Insurance." *Handbook of Health Economics* Volume 1, Chapter 11: 563-643.

Einav, Liran, Amy Finkelstein, and Paul Schrimpf. 2015. "The Response of Drug Expenditure to Nonlinear Contract Design: Evidence from Medicare Part D." *Quarterly Journal of Economics* 130(2): 841-899.

Einav, Liran, Amy Finkelstein, and Paul Schrimpf. 2016. "Bunching at the Kink: Implications for Spending Responses to Health Insurance Contracts." NBER Working Paper No. 22369.

Eliason, Paul, Paul Grieco, Ryan McDevitt, and James Roberts. 2016. "Strategic Patient Discharge". Unpublished mimeo, Duke University. Available at https://faculty.fuqua.duke.edu/~rcm26/LTCH.pdf.

Kim, Yan, Eric Kleerup, Patricia Ganz, Ninez Ponce, Karl Lorenz, and Jack Needleman. 2015. "Medicare Payment Policy Creates Incentives for Long-Term Care Hospitals to Time Discharges for Maximum Reimbursement." *Health Affairs* 34(6): 907-915.

Liu, Korbin, Cristina Baseggio, Douglas Wissoker, Stephanie Maxwell, Jennifer Haley, and Sharon Long. 2001. "Long-Term Care Hospitals Under Medicare: Facility- Level Characteristics." *Health Care Financing Review* 23(2): 1-18.

Manoli, Dayanand, and Andrea Weber. forthcoming. "Nonparametric Evidence on the Effects of Financial Incentives on Retirement Decisions." *American Economic Journal: Economic Policy.*

Medicare Payment Advisory Commission (MedPAC). 2004. "A Data Book: Health Care Spending and the Medicare Program."

Medicare Payment Advisory Commission (MedPAC). 2015a. "A Data Book: Health Care Spending and the Medicare Program."

Medicare Payment Advisory Commission (MedPAC). 2015b. "Report to the Congress: Medicare Payment Policy, Chapter 7: Medicare's Post-Acute Care: Trends and Ways to Rationalize Payments."

Newhouse, Joseph, Alan Garber, Robin Graham, Margaret McCoy, Michelle Mancher, and Ashna Kibria. 2013. "Variation in Health Care Spending: Target Decision Making, Not Geography." Institute of Medicine. Available at http://www.nationalacademies.org/hmd/~/media/Files/Report%20Files/2013/Geographic-Variation2/geovariation_rb.pdf

Tauchen, George. 1986. "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions." *Economics Letters* 20(2): 177–181.

Weaver, Christopher, Anna Wilde Mathews, and Tom McGinty. 2015. "Hospital Discharges Rise at Lucrative Times." *The Wall Street Journal*, February 17, 2015.

Appendix

Appendix A. LTCH Payments

TBA

Appendix B. Computation and Estimation

TBA

Appendix C. Post-Discharge Cost Projection

TBA

Figure 1: LTCH payment schedules before and after PPS

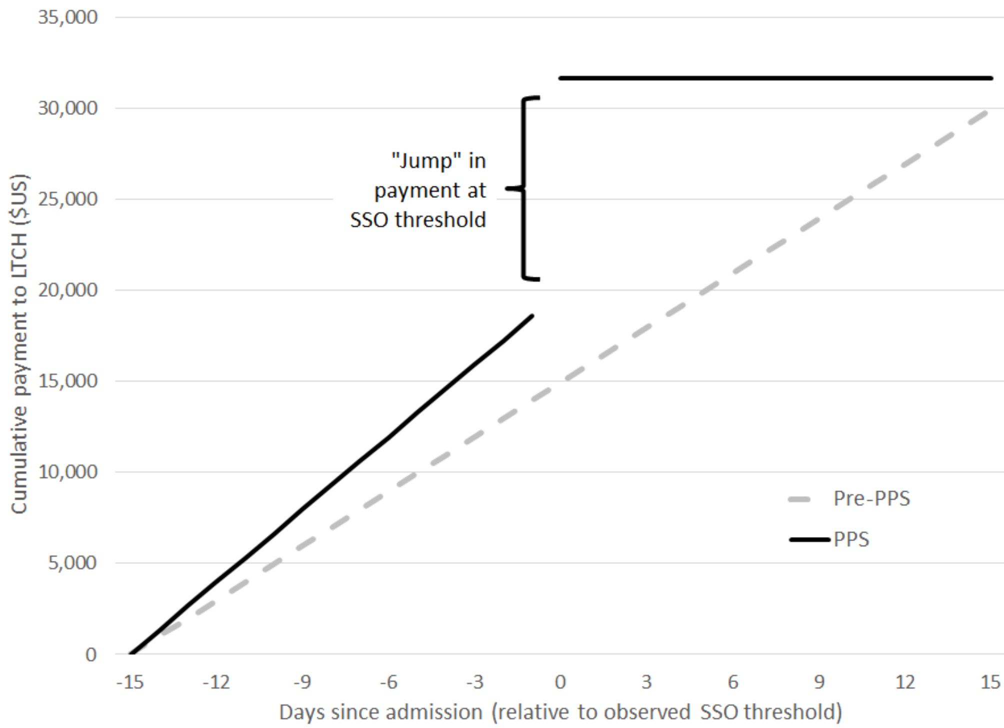

Figure presents the payment schedule (in 2012 dollars) in both the pre-PPS and PPS periods. Sample pools admissions that are associated with different SSO thresholds, and x-axis is normalized by counting days relative to the threshold. The linear payment schedule begins with the first day of admission, and the y-axis is normalized to zero for day -15.

Figure 2: Patents flow into and out of Post-Acute Care

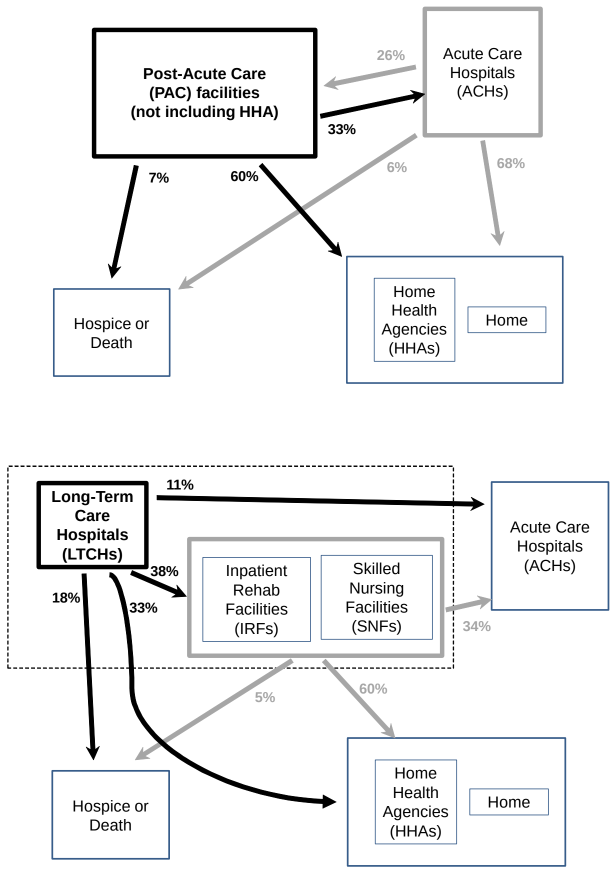

Top panel shows patient flow from acute care hospitals (ACHs) to the different destinations: post-acute care (PAC), home and home health agencies, and death or hospice. Bottom panel shows how the pattern is different, within PAC, between Long-Term Care Hospitals (LTCHs) and other PAC facilities (SNFs and IRFs). All numbers for this figure use the universe of Traditional Medicare admissions during the PPS period (Oct 2007 to Sept 2012). Numbers are shares of total discharges from each type of facility, excluding a small share of discharges (never greater than 5%) that are more difficult to classify. See Appendix for more details.

Figure 3: Discharge patterns by length of stay

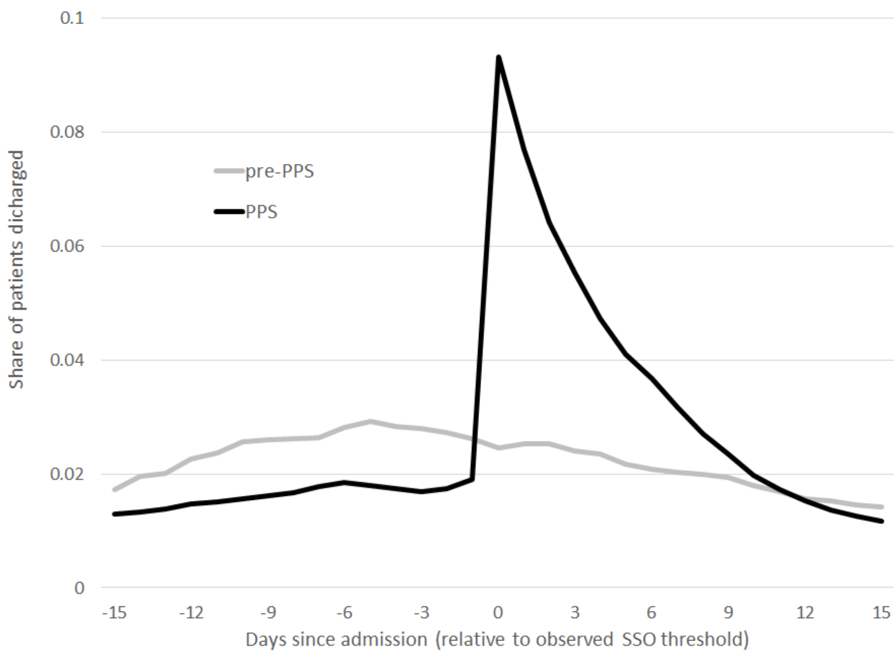

Figure presents the distribution of the time of discharge relative to the SSO threshold. That is, each number graphed represents the number of discharges at a given (relative) day divided by the total number of LTCH admissions. Sample pools admissions that are associated with different SSO thresholds, and x-axis is normalized by counting days relative to the threshold.

Figure 4: Discharge patterns across discharge destinations

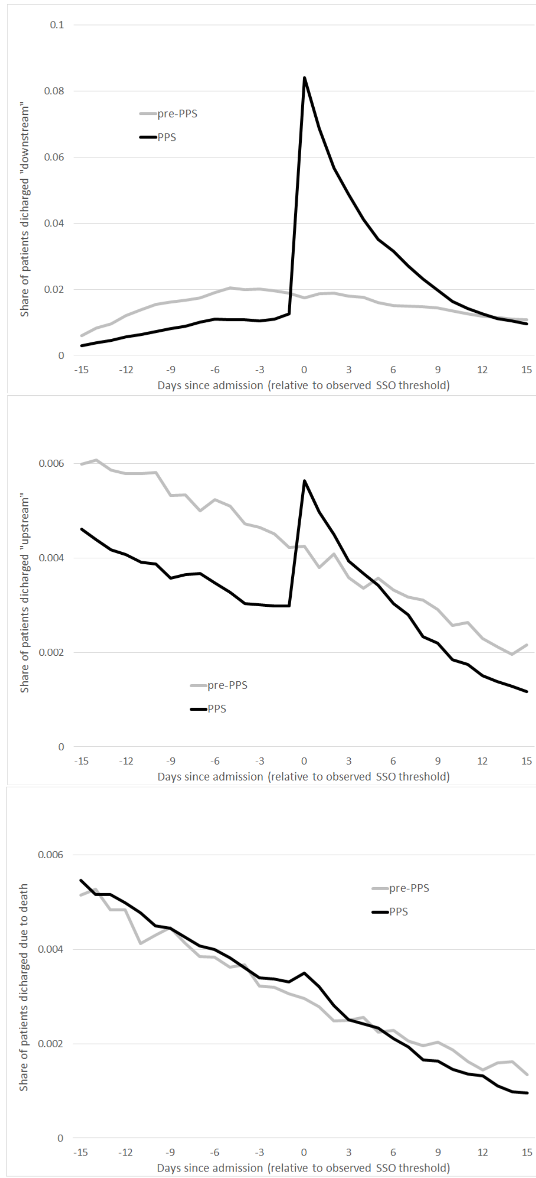

Figure is similar to Figure 3, but presents the distribution separately by discharge destination. Top panel presents discharges "downstream" (to SNF, IRF, or home), middle panel presents discharges "upstream" (ACH or hospice), and the bottom panel presents discharges due to death.

Figure 5: Post-discharge costs

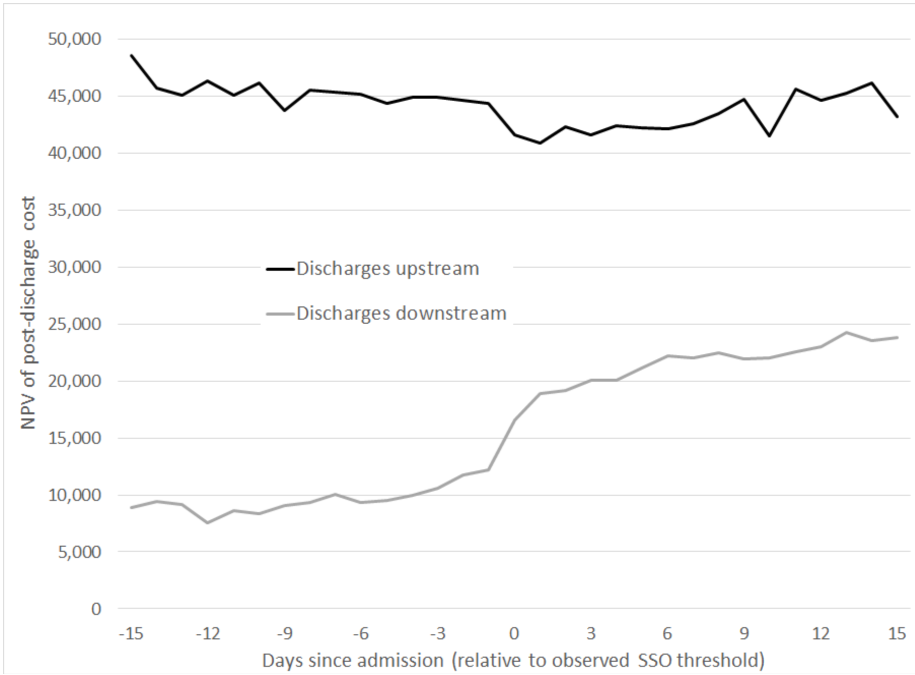

Figure presents the post-discharge costs, by discharge day and discharge destination. To construct it, we follow each patient discharged from the LTCH and add up total TM costs associated with the recovery episode. We then average across all discharges by day and destination. We define a recovery episode as ongoing until there is a break of at least two days that does not involve a facility stay.

Figure 6: Mortality patterns by length of stay

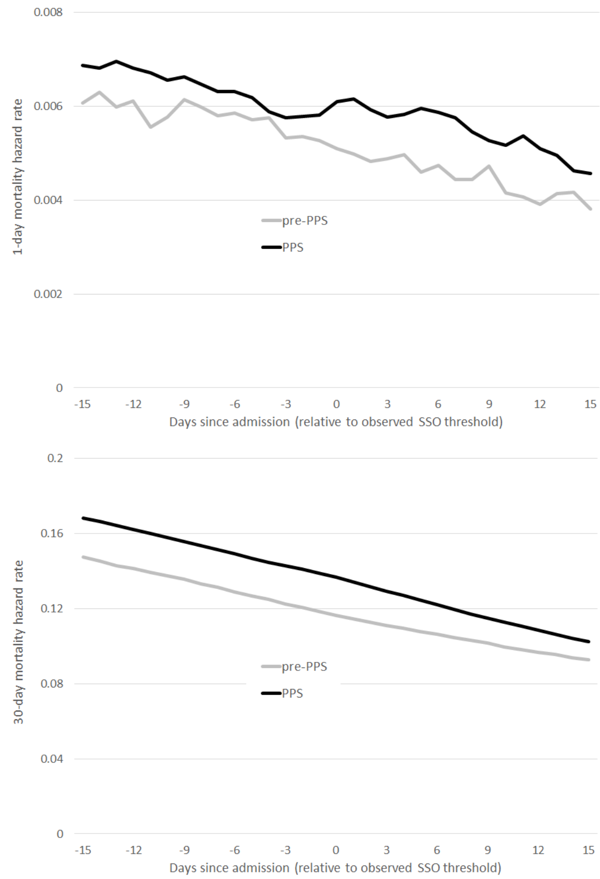

Figure presents mortality hazard rates by day. Mortality includes any mortality, whether it occurs within the LTCH or after discharge. Each panel presents hazard rates for different subsequent horizons: same day (top) and 30-day forward (bottom).

Figure 7: Differential patterns for for-profits and not-for-profit LTCHs

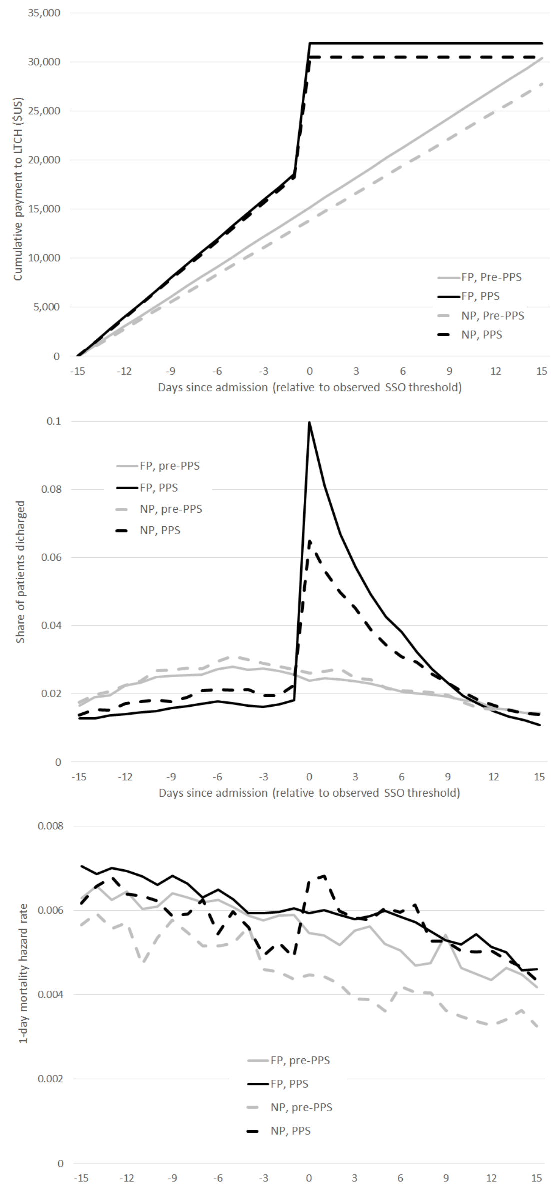

Figure replicates earlier figures, but separates the analysis for for-profit and non-profit LTCHs (the latter includes government-operated LTCHs). The top panel reports the payment schedule, replicating Figure 1. The middle panel reports discharge patterns, replicating Figure 3. The bottom panel reports mortality patterns, replicating the top panel of Figure 6.

Figure 8: Model fit

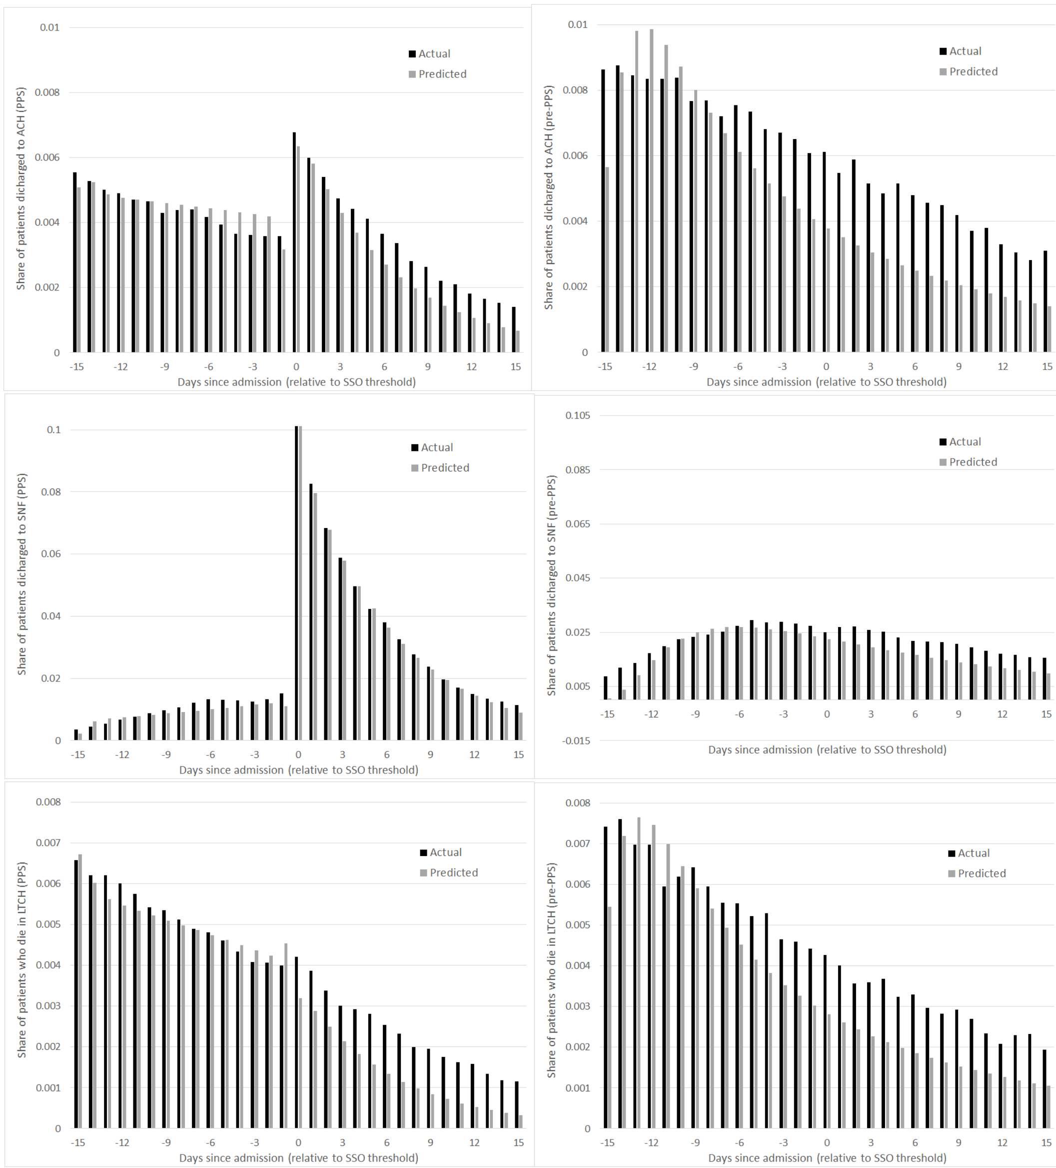

Figure shows the moments we use for estimation, and how the model is able to fit them. Black bars in each panel represent the actual moments from the data, and the gray bars represent the predicted moments from the model estimates. The left three panel represent the PPS period, and the right three panels represent the pre-PPS period. The top panels show discharge rates to ACH, the middle panels show discharge rates to SNF, and the bottom panels show mortality rates (within the LTCH).

Figure 9: Implied health processes and optimal discharge policy

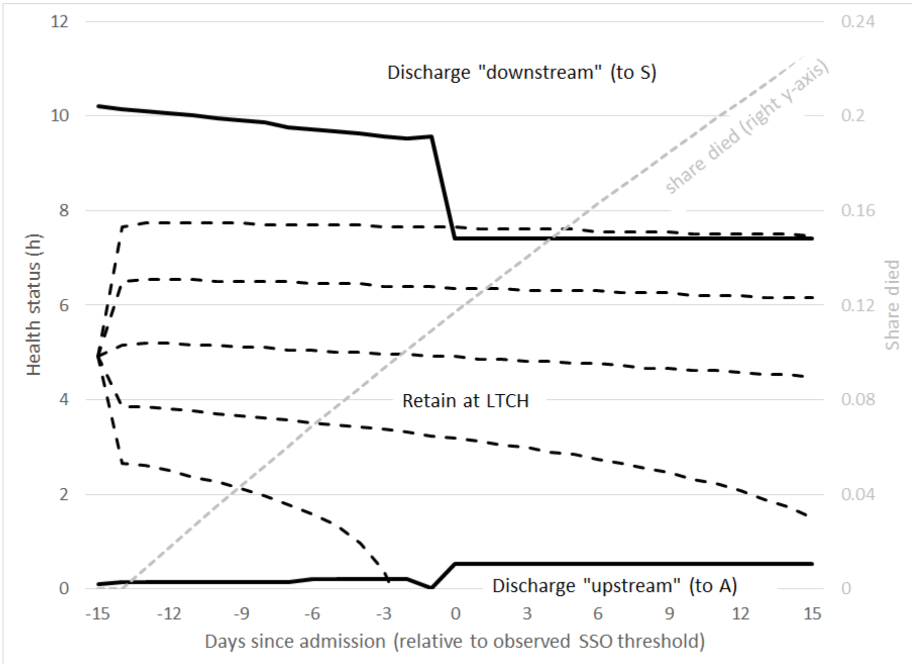

Figure describes the implications of the estimated model. The two black solid lines represent the policy function. The top black line approximates the health level above which a patient is discharged to S, and the bottom black line approximates the health level below which a patient is discharged to A. The black dashed lines show the health status distribution of patients who got admitted (and day -15) at a median health status, and then followed the estimated health process. Each line represents a different percentile (from top to bottom: 90th, 75th, median, 25th, and 10th). Recall that a health status of zero implies mortality, which is an absorbing state. The dashed gray line represents the cumulative share of individuals who died prior to that day (again, conditional on entering the LTCH with median health status).

Figure 10: Counterfactual payment schedules

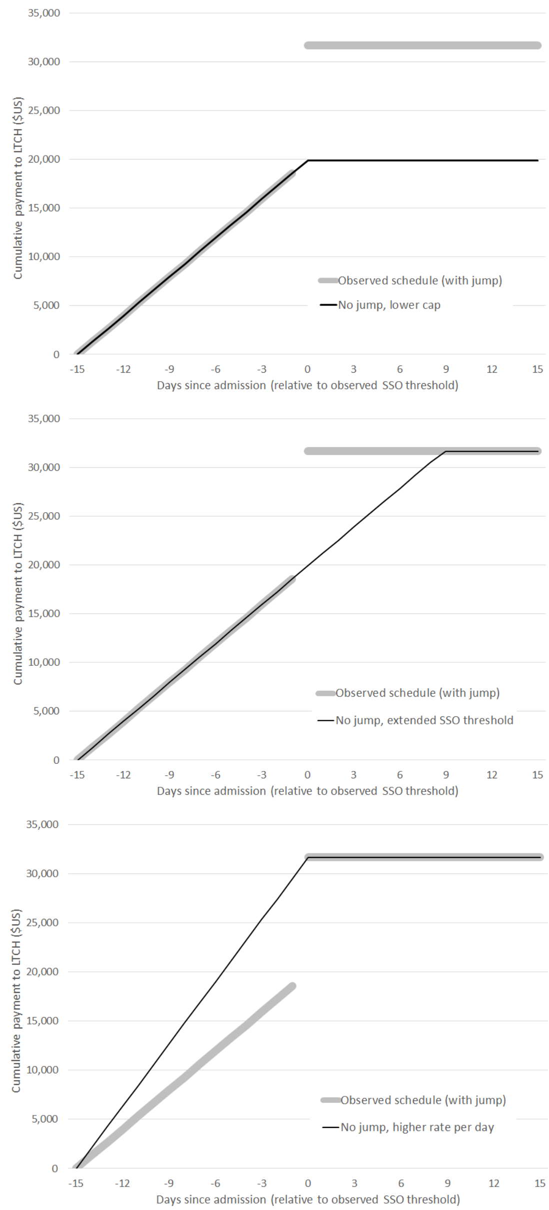

Figure shows the observed (PPS) payment schedule (thick gray line in all panels) and the three counterfactual payment schedules we consider (black line in each panel). All counterfactual schedules eliminate the jump in payments at the SSO threshold, but do this in different, increasingly generous (from top to bottom) ways.

Figure 11: Counterfactual discharge patterns

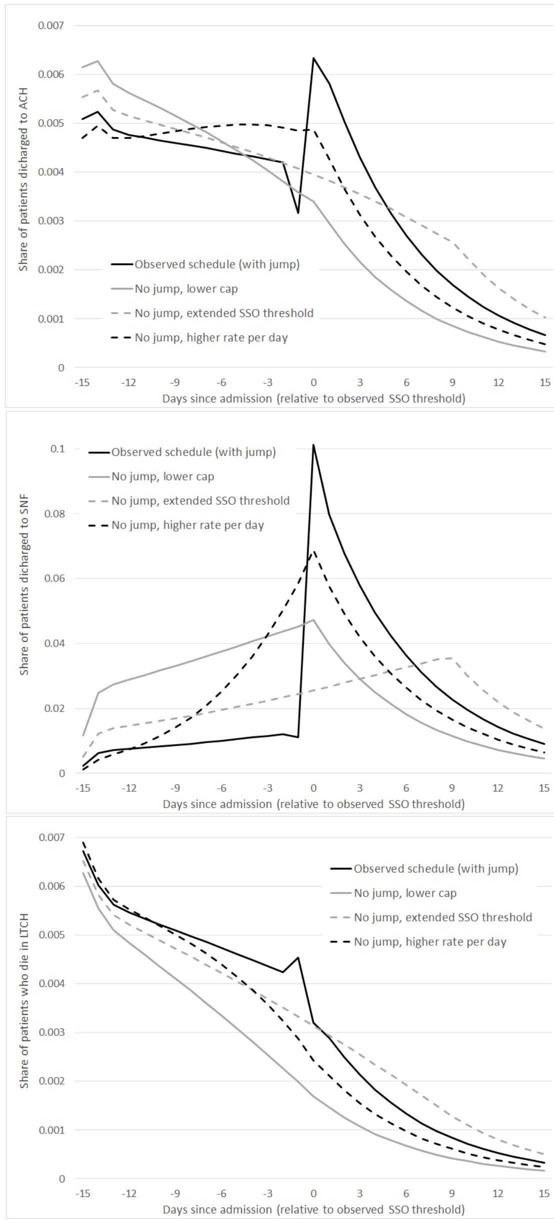

Figure show discharge and within-LTCH mortality patterns from three counterfactual payment schedules. The solid black line reports results that are based on our parameter estimates (reported in Table 2) and the observed payment schedule, and each other line reports the results from predicted discharge and mortality patterns under a different counterfactual payment schedule. The counterfactual payment schedules we consider are described in the main text and illustrated in Figure 10.

Table 1: Summary statistics

	Pre-PPS (Jan 2000 - Sep 2002)			PPS (Oct 2007 - Sep 2012)		
	ACH	LTCH	SNF/IRF	ACH	LTCH	SNF/IRF
Number of admissions (000s)	29,223	220	5,058	51,191	635	12,264
Panel A. Patient attributes						
Average age	74.5	73.9	80.2	73.5	71.7	79.1
Fraction male	0.43	0.44	0.35	0.44	0.48	0.37
Fraction white	0.84	0.74	0.87	0.82	0.73	0.85
Fraction black	0.11	0.20	0.10	0.12	0.20	0.11
Fraction qualified for Medicare via age	0.86	0.83	0.94	0.81	0.76	0.91
Fraction dual eligible	0.24	0.31	0.27	0.27	0.38	0.27
Panel B. Patient health indicators						
Number of Chronic Conditions	5.6	6.4	6.6	6.7	7.8	7.9
Number of inpatient days in prev. year						
30 Day Mortality Since Admission	0.082	0.149	0.114	0.079	0.160	0.087
90 Day Mortality Since Admission	0.142	0.283	0.219	0.139	0.310	0.185
Fraction who are home after 90 days						
Three most common DRGs:						
Panel C. Procedures during stay						
Fraction with no procedures	0.428	0.613	0.949	0.397	0.281	0.976
Number of procedures (cond. on any)	2.51	2.43	1.97	2.70	2.91	2.12
Three most common procedures:	Transfusion (6.2%) Arteriography (5.5%) Cardiac cath. (5.2%)	Cath (7.4%) Transfusion (5.5%) Occ. Therapy (4.9%)	Phys. Therapy (2.6%) Occ. Therapy (2.4%) Transfusion (0.3%)	Transfusion (10.3%) Cath. (6.5%) Dialysis (4.6%)	Cath (19.3%) Transfusion (18.4%) Ventilation (14.4%)	Occ. Therapy (1.2%) Phys. Therapy (1.2%) Transfusion (0.2%)
Length of stay	5.6	26.6	23.5	5.2	25.5	25.9
Panel D. Payments and cost						
Total Medicare payments per stay	9,425	28,386	9,727	10,865	35,402	13,124
Medicare payments per day	1,671	1,068	414	2,087	1,390	507
Out-of-pocket payments	772	2,344	1,568	839	1,916	1,990
Out-of-pocket payments per day	137	88	67	161	75	77
Total reported costs	--	28,386	--	--	36,218	--
Reported cost per day	--	1,068	--	--	1,422	--

Some of the entries in the table are empty. They will get populated in future versions of this paper.

Table 2: Parameter estimates

	Parameter	Std. error
Health process during pre-PPS:		
μ_0	8.32	
σ_0	3.46	
μ	0.42	
ρ	0.99	
σ	2.27	
Health process during PPS:		
μ_0	4.91	
σ_0	1.86	
μ	4.26	
ρ	0.19	
σ	1.95	
Preferences:		
γ	0.93	
β	0.9999	
ν_1^A (000s)	-35.69	
ν_0^S (000s)	-72.85	
ν_1^S (000s)	6.27	

Table presents parameter estimates of the parameters in our baseline specification. Standard errors will get populated in future versions of this paper.

Table 3: Discharges and payments from counterfactuals scenarios

	Observed sched. (1)	Lower cap (2)	Extended SSO (3)	Higher per day (4)
LTCH payments:				
Total payments	26,294	14,543	21,946	26,854
Average LOS*	18.2	13.5	18.5	16.4
Payment per day	1,444	1,079	1,188	1,642
Discharges to ACH:				
Total payments	5,104	4,332	5,359	4,762
Share of discharges	0.11	0.10	0.12	0.11
Payment per discharge	44,427	44,027	44,397	44,318
Discharges to SNF:				
Total payments	15,302	14,883	12,993	14,975
Share of discharges	0.79	0.83	0.78	0.80
Payment per discharge	19,485	17,903	16,657	18,627
Total Medicare payments	46,700	33,758	40,298	46,591

Table presents results from three counterfactual payment schedules. Column (1) reports results that are based on our parameter estimates (reported in Table 2) and the observed payment schedule, and each other column reports the results from predicted discharge patterns under a different counterfactual payment schedule. The counterfactual payment schedules we consider are described in the main text and illustrated in Figure 10.

* Length of stay is measured from day -15. To make it comparable to the summary statistics reported in Table 1, all numbers should be increased by 7.5 days (because the average SSO threshold across admissions in our sample is 22.5 days).

Table 4: Results from additional counterfactuals exercises

	Observed (PPS) (1)	All For-Profit (2)	All Non-Profit (3)	Pre-PPS admission mix (4)
LTCH payments:				
Total payments	26,294	26,671	24,943	
Average LOS*	18.2	18.2	17.9	
Payment per day	1,444	1,466	1,391	
Discharges to ACH:				
Total payments	5,104	4,906	4,969	
Share of discharges	0.11	0.11	0.11	
Payment per discharge	444,277	44,600	44,542	
Discharges to SNF:				
Total payments	15,302	15,440	15,031	
Share of discharges	0.79	0.79	0.79	
Payment per discharge	19,485	19,610	19,045	
Total Medicare payments	46,700	47,017	44,944	

Table reports the discharge and payment patterns under the observed schedule (column (1)), as well as two counterfactual scenarios, where we apply the model estimates from for-profit hospitals only and from non-profit hospitals only (see Appendix Table 2) to all LTCHs hospitals. In future version of this paper, we will also report (in column (4), which is currently empty) a third counterfactual, where we use the pre-PPS health process and the PPS payment schedule in order to quantify the importance of the change in the patient mix.

* Length of stay is measured from day -15. To make it comparable to the summary statistics reported in Table 1, all numbers should be increased by 7.5 days (because the average SSO threshold across admissions in our sample is 22.5 days).

Appendix Figure 1: Discharge patterns, re-weighted

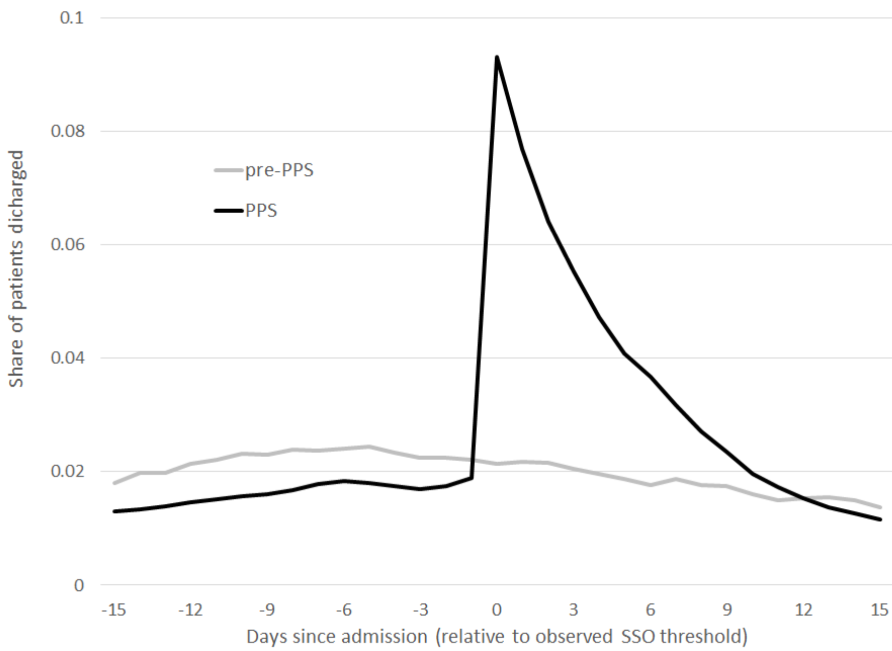

Figure is the same as Figure 3 in the main text, except that pre-PPS line is re-weighted to reflect the same DRG mix as in the PPS period.

Appendix Figure 2: Post-discharge mortality rates

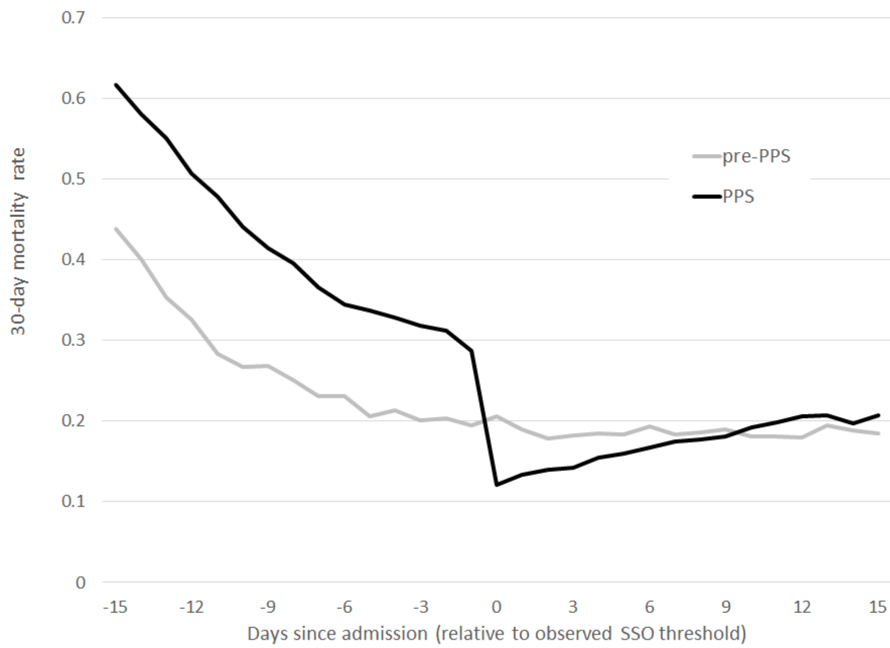

Figure presents the (forward looking) 30-day mortality rate after discharge date as a function of the day of discharge.

Appendix Figure 3: Discharge patterns, by for-profits status

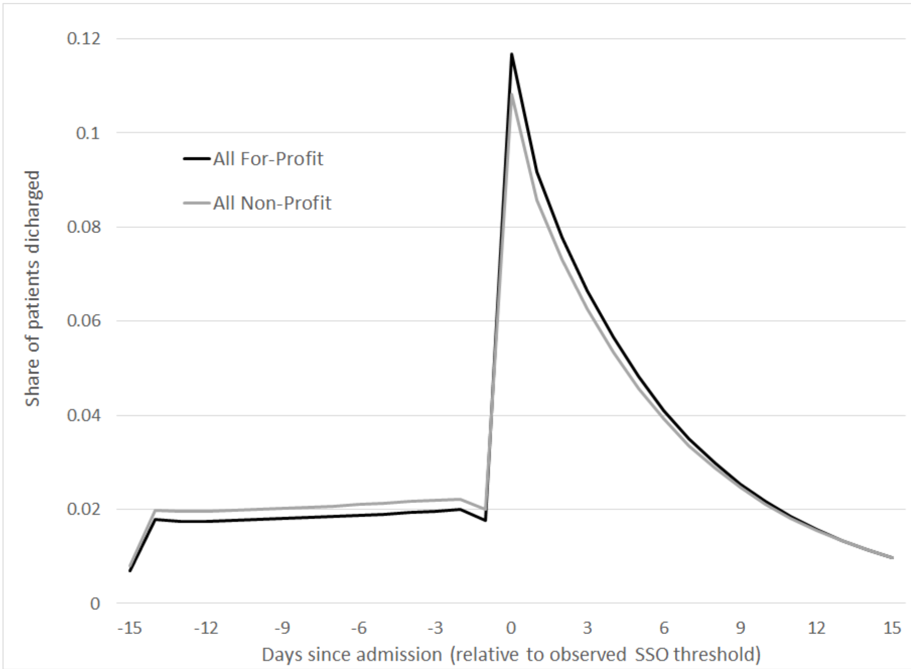

Figure presents the counterfactual discharge patterns associated with the counterfactuals reported in columns (2) and (3) of Table 4.

Appendix Table 1: Parameter estimates, by for-profits status

	For-profit LTCHs		Non-profit and gov-owned LTCHs	
	Parameter	Std. error	Parameter	Std. error
Health process during pre-PPS:				
μ_0	7.97		7.53	
σ_0	3.16		3.17	
μ	0.41		0.43	
ρ	0.99		0.99	
σ	2.20		2.12	
Health process during PPS:				
μ_0	4.89		5.11	
σ_0	1.85		1.98	
μ	4.23		4.67	
ρ	0.19		0.16	
σ	1.93		2.14	
Preferences:				
γ	0.88		0.92	
β	0.9999		0.9999	
v_1^A (000s)	-27.37		-37.53	
v_0^S (000s)	-65.89		-83.56	
v_1^S (000s)	6.14		7.43	

Table presents estimates of the parameters, when we estimate the model separately for for-profit and non-profit hospitals. Standard errors will get populated in future versions of this paper.