MANAGEMENT SCIENCE
Vol. 59, No. 5, May 2013, pp. 1027–1044
ISSN 0025-1909 (print) | ISSN 1526-5501 (online)
http://dx.doi.org/10.1287/mnsc.1120.1606

On Hospice Operations Under Medicare Reimbursement Policies

Barış Ata
Northwestern University, Evanston, Illinois 60208, b-ata@kellogg.northwestern.edu

Bradley L. Killaly
Emory University, Atlanta, Georgia 30322, brad.killaly@emory.edu

Tava Lennon Olsen
University of Auckland, Auckland 1142, New Zealand, t.olsen@auckland.ac.nz

Rodney P. Parker
University of Chicago, Chicago, Illinois 60637, rodney.parker@chicagobooth.edu

This paper analyzes the United States Medicare hospice reimbursement policy. The existing policy consists of a daily payment for each patient under care with a global cap of revenues accrued during the Medicare year, which increases with each newly admitted patient. We investigate the hospice's expected profit and provide reasons for a spate of recent provider bankruptcies related to the reimbursement policy; recommendations to alleviate these problems are given. We also analyze a hospice's incentives for patient management, finding several unintended consequences of the Medicare reimbursement policy. Specifically, a hospice may seek short-lived patients (such as cancer patients) over patients with longer expected lengths of stay. The effort with which hospices seek out, or recruit, such patients will vary during the year. Furthermore, the effort they apply to actively discharge a patient whose condition has stabilized may also depend on the time of year. These phenomena are unintended and undesirable but are a direct consequence of the Medicare reimbursement policy. We propose an alternative reimbursement policy to ameliorate these shortcomings.

Key words: healthcare; government: regulations; fluid analysis; simulation
History: Received September 27, 2010; accepted June 5, 2012, by Christian Terwiesch, operations management. Published online in *Articles in Advance* December 10, 2012.

1. Introduction and Problem Motivation

Hospices are healthcare providers that cater to patients in the end phases of their lives who choose to undertake palliative care in lieu of further curative care. Since 1983, Medicare has reimbursed hospices for the care provided to eligible patients. This reimbursement consists of a daily payment for care, but the American federal government's exposure is limited by an annual cap that is dependent on the number of patients admitted in a year. The Medicare hospice benefit is generally regarded as a success because it improves the quality of life while saving Medicare (compared with a patient continuing curative care) an average of $2,309 per hospice user (Taylor et al. 2007). However, some recent issues have arisen to question the efficacy of this reimbursement policy. Specifically, an increasing number of providers have declared bankruptcy, and the blame for this is attributed to the structure of the reimbursement policy (Sack 2007). We investigate this claim analytically, study hospice provider incentives under the current policy, and explore alternative possibilities to the current Medicare policy.

More specifically, we formulate a model for hospice profit and use it to examine the potential causes for hospices receiving payments that exceed the cap and the reasons behind potential bankruptcies. However, the payment scheme elapsing over a finite horizon raises further issues beyond the profitability of the provider. The annualized accounting involved may be leading to some undesirable traits in the rate of hospice admissions and discharges, such as patient recruiting and discharge rates that differ across diseases and change during the year. If this is the case, it is likely contrary to the United States government's equity objectives.

There is indeed evidence to suggest that some untoward patient management practices occur in the hospice industry, some of which are compelled by the reimbursement policy's cap. Jenkins et al. (2011) report hospice providers confirm that they modify their practices when they are in danger of exceeding their cap. In their survey of 55% of the hospice providers in Alabama (a state frequently mentioned by the Medicare Payment Advisory Commission (MedPAC) as having providers that exceed their cap; see MedPAC 2010, for example), 24.4% of respondees reported modifying their actions when faced with the prospect of a binding cap. Of such modifications, the two most cited behaviors were discharging patients (17% of all respondees) and marketing "to a certain type of patient (e.g., cancer patients)" (10.4% of all respondees) to alleviate cap problems. We seek to model these two behaviors in our work. Actions such as these are a sufficient concern such that MedPAC has advised the Secretary of Health and Human Services to direct the Office of the Inspector General to investigate the financial relationships between "hospices and long-term care facilities such as nursing facilities and assisted living facilities that may represent a conflict of interest and influence admissions to hospice" and "the appropriateness of hospice marketing materials and other admissions practices and potential correlations between length of stay and deficiencies in marketing or admissions practices" (MedPAC 2010, p. 147).

This paper investigates the hospice manager's optimal recruitment problem and finds that, indeed, the manager has an incentive to purposefully seek cancer and other short-lived patients when the hospice's cap might be exceeded. We also investigate the incentives behind the practice of "live discharges" (MedPAC 2010). A *live discharge* is a living patient who is released from a hospice. Given the unpredictable trajectory of terminal diseases, some patients' diseases simply do not follow expectations. A patient may recover, or a patient may simply elect to resume conventional medical treatment, which of course results in their leaving the hospice. However, a hospice caring for patients who have exceeded their contribution to the cap may feel the pressure to discharge such patients, however unethical (and potentially illegal) such a practice may be, a possibility we explicitly include. Our model shows that the hospice manager may find it optimal to live-discharge patients particularly toward the end of the fiscal year if the current cap position is not desirable. Moreover, in those cases the rate of live-discharges increases toward the end of the fiscal year. To remedy these unintended consequences, we propose an alternative reimbursement policy and show that it indeed alleviates these nonstationary recruiting or live-discharge patterns.

The remainder of the paper is structured as follows. Section 2 briefly describes the existing Medicare reimbursement policy and presents the institutional details and motivation for our work, including Medicare and hospice-centered references. In §3 the hospice profitability model is formulated and analyzed, and remedies to the policy's shortcomings are considered. In §4 we formulate and analyze a dynamic deterministic recruitment and discharge model, highlighting several disturbing unintended consequences of the existing policy. We also perform a simulation study of the effects of stochasticity on our conclusions. Section 5 presents a policy to overcome the unintended consequences of the existing policy. Concluding remarks are provided in §6. Proofs are provided in Online Appendices A and B; Online Appendix C contains details of our simulation study. These online appendices are provided at http://faculty.chicagobooth.edu/rodney.parker/research/hospicesappendix.pdf.

2. Medicare Reimbursement Policy and Literature Review

We begin with a review of the Medicare reimbursement policy features that are relevant for our purposes. The reader is referred to MedPAC (2010) for a more comprehensive description of the current policy. We then review the relevant literature on the hospice system and Medicare reimbursement program.

The Medicare year runs from November 1 through October 31 of the following year. To be admitted to a hospice, a patient needs the signature of two physicians (typically, one is the patient's primary attending physician and the other is employed by the hospice), certifying that the patient is not expected to live more than six months from admission[1] and that the patient agrees to forgo any curative care and undertake palliative care only.

During the Medicare year, a hospice receives a payment for a patient under hospice care for a part of or an entire day. This payment differs according to whether the patient is receiving routine home care ($142.91 per day), continuous home care ($834.10 per day), inpatient respite care ($147.83 per day), or general inpatient care ($635.74 per day).[2] These payment rates can differ slightly by region in the United States depending on estimates of the cost of operation in these regions, but they do not depend on the disease afflicting the patient. In 2002 and 2003, 93% of reimbursed hospice days were paid at the routine home care rate, 4.1% were continuous home care days, 2.7% were inpatient respite care days, and 0.2% were general inpatient care days (MedPAC 2006). We focus on routine home care in our models, as the vast majority of reimbursed days are for routine home care.

[1] If the patient lives beyond the initial six-month period, the patient can be recertified for two sequential 90-day periods, followed by an unlimited number of sequential 60-day periods. Until 1990, there was a limit of 210 days over which a hospice could receive payments for a patient.

[2] These payment rates apply to the federal fiscal year 2009 (see http://www.cms.gov/Outreach-and-Education/Medicare-Learning-Network-MLN/MLNProducts/downloads/hospice_pay_sys_fs.pdf, accessed July 2009).

The second part of the Medicare hospice reimbursement policy is a payment cap, applied to the entire hospice (i.e., the cap is not patient specific),[3] intended to limit the government's exposure. At the beginning of the Medicare year, this cap is zero, but it increases by \$23,014.50[4] for every newly admitted patient. The daily payment rates and cap increase quantities are adjusted from year to year, but unlike the daily rates, the cap does not vary by geography. At the end of the Medicare year, if the hospice's cap is less than the total of the daily payments received, this excess amount must be repaid to Medicare. If the cap is greater than the payments received, then the hospice did not receive as many payments as they were eligible for but no adjustment is made. The cap is reset to zero at the beginning of the new Medicare year.

Since the creation of the Medicare hospice benefit in 1983, there has been much research on many aspects of hospices, but, to the best of our knowledge, none has directly addressed the issues focused on here. In fact, much of the literature does not address the reimbursement policy at all. An exception is Fraser (1985), who describes the ethical and policy implications of Medicare's hospice reimbursement policy, but does not identify the issues of provider bankruptcy or nonstationary and disease-dependent recruitment and discharge, the foci of our study. The U.S. Government Accountability Office (2004) investigates whether modifications to the reimbursement policy are warranted but limits its focus to the comparison of the per-diem rates and the costs of care. Huskamp et al. (2001) consider how the Medicare rules affect care. There is consistent evidence in the literature that the hospice benefit reduces Medicare costs (e.g., Campbell et al. 2004) while enhancing end-of-life care (Pyenson et al. 2004, Taylor et al. 2007).

Previous empirical work examines common characteristics of hospices that exceed their caps. In particular, MedPAC (2010) reports the following attributes: they tend to be for-profit, freestanding facilities;[5] have a smaller patient census; treat a larger share of patients with Alzheimer's disease and other neurological conditions; exhibit significantly longer lengths of stay (LOS), even when patient mix is taken into account; and have a proportion of patients with stays exceeding 180 days (for particular diseases) substantially higher than those hospices below the cap.

[3] Medicare applies a second cap, which is rarely enacted. It limits the proportion of inpatient care days to 20% of all reimbursed days. This is to encourage hospice care in the patient's home as the primary method of delivery. Any days exceeding the 20% level will be reimbursed at the routine home care rate.

[4] Again, for the year 2009.

[5] Freestanding facilities are those not operated by a hospital, home health agency, or skilled nursing facility.

Furthermore, there is an increasing trend of hospices receiving payments greater than what they were permitted (i.e., exceeding their caps) and having to repay this excess to the federal government. For example, MedPAC (2010) reports the percentage of all hospices exceeding their caps rose from 2.6% in 2002 to 10.4% in 2007 and points to two explanatory factors. First, there has been an increase in the proportion of longer-stay patients. Second, there has been an increase in the LOS of the longer-stay patients. In 1998, 47% of all hospice users had noncancer diagnoses; this had risen to 69% by 2008 (MedPAC 2010). MedPAC (2010) reports the median LOS remained steady at 17 days between 2000 and 2008, but the 90th percentile grew from 141 days to 235 days. In brief, the short stays remained at a similar LOS, but the long stays grew longer. There appear to be no definitive explanations for these extended LOS, although MedPAC was sufficiently concerned that they recommended Congress direct the Secretary of Health and Human Services to require a medical review of all stays exceeding 180 days in hospices where such stays make up 40% or more of all cases (MedPAC 2010). If the trends of increasing proportions of noncancer patients and longer life spans for the longest-living noncancer patients continue, we envision there will be more and more financially distressed hospices, particularly if they happen to be smaller providers. The results in our paper suggest this will lead to further and more extreme distortions of incentives. Given that MedPAC is aware of these trends, we suspect that they will intervene to preserve the hospice benefit for truly terminal patients, either through stricter enforcement of the six-month LOS estimate or through another mechanism. Modeling this is beyond the scope of our paper.

Concerning hospice costs, MedPAC (2010) finds that the average provider costs per day can vary by hospice type, that for-profit-based hospices are less costly than nonprofit hospices, that rural hospices are less costly than urban hospices, and, curiously, that hospices exceeding the payment cap are less costly than those below the cap. It also found that the daily costs are higher at admission and discharge than regular care, so providers with longer average LOS have lower daily costs, which may explain why hospices exceeding the payment cap are less costly. Killaly and Mukamel (2010) find that the providers' marginal and average costs are higher for cancer patients than noncancer patients. MedPAC also examines hospice margins; generally, all hospice types are profitable except hospital-based providers (which had consistently negative margins from 2001 to 2007, presumably the result of greater overhead allocation) and providers with the lowest patient volumes. MedPAC (2010) finds that hospice margins increase with patient volume in every year under study but projects the aggregate margin to drop from 5.9% in 2007 to 4.6% in 2010 across all hospices.

As to live discharges, Taylor et al. (2008) suggest that 15.5% of hospice users were discharged alive from the years 1993 to 2000. MedPAC (2010) finds that live discharges are far more prevalent amongst providers exceeding their caps (46% of all discharges) than those not exceeding their caps (16%) in 2007, and consequently recommended the Office of the Inspector General investigate the "appropriateness of enrollment practices for hospices with unusual utilization patterns" (p. 147). Carlson et al. (2009) find that live discharges are more prevalent for smaller hospices and that long-stay patients may be more susceptible to this practice, two factors that our findings suggest can lead to diminished profitability, although newer hospices were also commonly found to be live dischargers, perhaps implying that some inexperience in judging LOS may be a factor.

3. Model of the Current Policy

This section presents a static model of the existing Medicare policy for hospice reimbursement. We consider elements of the industry and market that may affect hospice profitability, including patient census, patient disease mix, and LOS uncertainty. The detrimental nature of poor mix realization, lack of scale, and uncertainty are well recognized in the operations management literature. For example, Eppen (1979) recognizes the value of pooling inventory in the context of warehouses. Such lessons are instructive for analyzing hospice operations, and we create a model to do so. To the best of our knowledge, there does not appear to be work in the operations management literature dealing with a cap akin to the Medicare hospice cap, although there are inventory papers with shared limited production capacity (e.g., Evans 1967). Recently, there has been a substantial increase in the interest of applying operations management techniques to healthcare-related topics (see, for example, Brandeau et al. 2004).

We express the general form for the hospice's profit given the Medicare reimbursement policy and then simplify the demand model to make it tractable. The fiscal year is considered as a whole; specific dynamics caused by year-end effects will be discussed in §4. In particular, all patients are assumed to arrive, be cared for, and be reimbursed for in the same fiscal year. Although in practice there will be patients that live from one Medicare year to the next, we feel that this model captures the key effects regarding hospice profitability at a high level without losing tractability or insights.

A hospice provider faces uncertainty concerning the patients' LOS. It is recognized that these uncertain LOS differ by patient disease (Christakis and Escarce 1996), although the Medicare reimbursement policy is independent of disease type. Our models consider only two disease types, although the results are robust to this characterization. Patients are classified as type 1 or type 2, characterized by the former having shorter mean LOS and the hospice incurring a lower marginal cost for the latter. In a broad sense, one could think of type 1 patients as those suffering from cancer and type 2 as those with noncancerous diseases, although this categorization is far from perfect because there are several noncancerous diseases with shorter LOS than several cancer diseases (e.g., chronic kidney disease has a mean LOS of 28–30 days; see Centers for Medicare and Medicaid Services 2009). For the remainder of the paper, we refer to two patient types.

Let N_1 be the number of type 1 (short LOS) patients admitted in a year and N_2 be the number of type 2 patients admitted. Assume that X_1^j is the remaining life span of type 1 patient j ($j = 1, \ldots, N_1$) and X_2^k is the remaining life span of type 2 patient k ($k = 1, \ldots, N_2$). Furthermore, assume type i patients cost c_i per day to treat, $i = 1, 2$, and let A be the fixed cost of operating the hospice. Let r be the daily precap reimbursement rate for patients and K the cap adjustment per admitted patient. Then

$$\text{Profit rate (per year)} = r\left(\sum_{j=1}^{N_1} X_1^j + \sum_{k=1}^{N_2} X_2^k\right) \wedge K(N_1 + N_2) - c_1 \sum_{j=1}^{N_1} X_1^j - c_2 \sum_{k=1}^{N_2} X_2^k - A.$$

As discussed above, this ignores beginning and end-of-horizon effects. We will refer to this model as our *static* model and relax this assumption with the *dynamic* model in §4.

For simplicity, we assume N_1 and N_2 are not random: $N_1 = \lambda_1$ and $N_2 = \lambda_2$. We also assume $\{X_1^j\}$ and $\{X_2^k\}$ are independent and identically distributed (i.i.d.) sequences and are independent of each other. Then, assuming λ_1 and λ_2 are sufficiently large, we conclude by the central limit theorem that

$$\text{Uncapped revenues} = r\left(\sum_{j=1}^{\lambda_1} X_1^j + \sum_{k=1}^{\lambda_2} X_2^k\right) \sim r\mathrm{N}(m, \sigma^2),$$

where $m = (\lambda_1 m_1 + \lambda_2 m_2)$, $\sigma^2 = (\lambda_1 \sigma_1^2 + \lambda_2 \sigma_2^2)$, and m_i, σ_i^2 denote the mean and the variance of X_i^1, respectively, for $i = 1, 2$. Let $\phi(\cdot)$ and $\Phi(\cdot)$ denote the probability density function and cumulative distribution function of the standard normal distribution, respectively. The following proposition characterizes the hospice's annual profit.

Proposition 1. *The expected annual* (*static*) *profit is given by*

$$\pi = \lambda\left\{K - (K - r\bar{m})\Phi\left(\frac{K/r - \bar{m}}{\bar{\sigma}/\sqrt{\lambda}}\right) - \frac{r\bar{\sigma}}{\sqrt{\lambda}}\phi\left(\frac{K/r - \bar{m}}{\bar{\sigma}/\sqrt{\lambda}}\right) - \bar{c} - \frac{A}{\lambda}\right\}, \qquad (1)$$

where $\lambda = \lambda_1 + \lambda_2$, $\bar{m} = (\lambda_1/\lambda)m_1 + (\lambda_2/\lambda)m_2$, $\bar{\sigma} = \sqrt{(\lambda_1/\lambda)\sigma_1^2 + (\lambda_2/\lambda)\sigma_2^2}$ *and* $\bar{c} = (\lambda_1/\lambda)(c_1 m_1) + (\lambda_2/\lambda)(c_2 m_2)$. *Static profit* π *is concave increasing in* K, *concave in* r, *decreasing in* σ_1 *and* σ_2, *and linearly decreasing in* c_1, c_2, *and* A.

The following insights are immediate from Proposition 1. First, high costs or low revenues obviously affect profits adversely. Second, so long as the per-patient margin is positive, large volumes are better for the hospice manager. Third, the hospice's profit decreases as the variability ($\bar{\sigma}$) increases. Finally, the effects of patient mix and LOS are also important in determining profitability.[6]

As mentioned in §1, hospices can go bankrupt for several reasons. The most important among these are volume and mix. A hospice with insufficient volume cannot cover its fixed costs. Furthermore, the negative impact of LOS variability diminishes as λ increases. To be specific, it can be seen from (1) that a hospice's profit loss due to LOS variability decreases as λ increases as a result of the term $\bar{\sigma}/\sqrt{\lambda}$, which captures effective LOS variability.

The second important reason for bankruptcy is improper mix. Intuitively, the extremes of the mix spectrum hurt the hospice profit. To see this, consider a hospice serving primarily cancer (short LOS) patients. Because its patients do not live very long, the hospice's per-patient revenue is low and, in particular, insufficient to cover its fixed cost unless it serves a very large number of patients. At the other extreme, consider a hospice serving primarily noncancer (long LOS) patients. Because its patients live a long time, the cap constraint will bind. Therefore, the hospice will accrue revenues during only a portion of a patient's stay (because the cap binds) whereas it incurs caring costs throughout the patient's stay, which may cause the hospice to go bankrupt. The ideal mix values are those that allow the hospice manager to leverage the benefits of both types of patients: cancer (short LOS) patients help build a large cap, but the hospice manager cannot take advantage of that cap with cancer patients alone, whereas noncancer (long LOS) patients help convert the cap to revenue.[7]

One remedy we propose to prevent possible bankruptcy is for appropriate hospices to merge. There are several potential benefits from this. First, the merged hospice will have larger volumes and expected profits are increasing in scale. Second, the merged hospice enjoys the well-known benefits of pooling (see, e.g., Eppen 1979) where the relative variability is reduced, and Proposition 1 shows that the hospice profit decreases in variability. Finally, and most important, if the constituent hospices are chosen well, the resulting patient mix in the merged provider could result in a more robust operating mix. For example, MedPAC (2010) highlights the issues of hospices in Alabama and Mississippi exceeding their caps, primarily as a result of the exceedingly long LOS; Sack (2007) highlights the dramatically shorter LOS in South Dakota. Such hospices could be merged.[8]

Despite merging hospices being attractive from a profit perspective, the practicalities of implementing such mergers must also be examined, which is beyond the scope of this paper. We will note that the government is in the prime position to act as a "matchmaker" to identify candidates for and encourage such mergers and, perhaps more practically, to remove any regulatory hurdles to inhibit the mergers. For example, hospices in different states with a common owner are not currently permitted to operate with a common cap. However, by doing so, they may also make the hospice benefit more expensive to Medicare (although still within their budgeted cap for admitted patients) and might need to reduce the hospice per diem to make such a recommendation revenue neutral, which is not without its own complications.[9] Further study on the precise cost implications may be necessary before implementation.

[6] Even though §4 shows that the annualized cap provides incentives for undesirable behavior, pooling the cap across the patients helps the hospice manager. If the cap were simply per patient, then the presence of variability in LOS could be extremely detrimental to the hospice because of a lack of risk pooling.

[7] To glean further insights from Proposition 1, consider the asymptotic regime where λ gets large while the mix remains the same. Then, using the fact that $\lim_{z\to\infty} z(1-\Phi(z))/\phi(z) = 1$ (see Zipkin 2000), the hospice's profit can be approximated for sufficiently large λ, as follows. If $K \neq r\bar{m}$, then $\pi \approx \lambda[K \wedge (r\bar{m}) - \bar{c} - A/\lambda]$. Similarly, if $K = r\bar{m}$, then $\pi \approx \lambda[K \wedge (r\bar{m}) - (r\bar{\sigma}\phi(0))/\sqrt{\lambda} - \bar{c} - A/\lambda]$. These equations highlight the importance of the mix and the cap through the term $K \wedge (r\bar{m}) - \bar{c}$, which shows that extreme values of mix may hurt profitability. These equations also validate that the hospice becomes more profitable as λ increases so long as per-patient margin is positive. The LOS variability has a first-order effect in profitability only when $K = r\bar{m}$, in which case increasing λ helps increase profit through risk pooling as well because of the term $\bar{\sigma}/\sqrt{\lambda}$. Otherwise, i.e., $K \neq r\bar{m}$, the revenue loss due to the imbalance between the cap and per-patient revenue dominate the revenue loss due to LOS variability.

[8] Mergers to correct mix imbalance only make sense if the patient mix is relatively stable across facilities.

[9] Medicare has sometimes reduced the hospice benefit per diem for some types of care.

As noted earlier, when carefully chosen, such mergers could potentially alleviate unfortunate systemic patient mix issues. Furthermore, we feel that correcting unprofitable patient mix through hospice mergers is a more socially equitable (and ethical) solution than allowing hospices to adjust their mix by actively searching for patients of the "right" type (inducing nonterminally ill patents to join the hospice) or live discharging patients (who are eligible for the benefit) to improve the hospice's position relative to the cap; under this latter solution, all patient types may no longer receive equal hospice access. In the following section, we show that hospices have an incentive to actively search for the right patient type and to discharge certain living patients to improve their position relative to the cap, and that the optimal search and live-discharge strategies change over the course of a year, further exacerbating equal access concerns.

4. The Hospice Manager's Problem

This section presents the patient management challenge a hospice manager faces to preserve profitability. We formulate and analyze a dynamic deterministic model where there are regular arrivals of hospice patients of two types during the Medicare year. In addition to these regular arrivals (at differing rates) of patients, hospices are allowed to seek, or "recruit," additional admissions of these patient types during the year and are allowed to discharge living patients whose conditions have stabilized. The recruiting rates of additional patients (of each type) and the discharge rates for the pool of stable patients are the decision variables of this dynamic model. Our model assumes that hospices have the ability to recruit beyond their natural arrival rates but that such recruiting has a convex increasing cost because diseconomies of scale appear natural in a limited and competitive market. Similarly, we also allow the hospice manager to discharge living patients should their conditions stabilize, also with a convex increasing cost. Note that we do this in an effort to inform policy makers of the incentives inherent in the current system; we do not propose these rates as a prescription for hospice management.

4.1. A Fluid Model

The model advanced in this section treats arrivals of patients to each class as fluid, arriving at the system at a constant rate. In particular, class i customers arrive at rate λ_i $(i=1,2)$ per time unit. Although all patients are diagnosed as terminally ill at admission to the hospice, the diagnosis may turn out to be false in some cases, and those patients may be discharged alive (should their conditions stabilize). Let a and b index the patients who are truly terminally ill and those who may be discharged alive, respectively. For $i=1,2$, γ_i^a and γ_i^b denote their fraction $(\gamma_i^a+\gamma_i^b=1)$. We take this approach because assuming all patients may be discharged would be problematic, as discharging the truly terminal is both immoral and illegal. Our model is based on a single fiscal year, with time indexed by $t\in[0,T]$.

Terminally ill patients have different LOS distributions than patients who are misdiagnosed. The life span of a class i,j patient has mean m_i^j and density $f_i^j(\cdot)$ for $i=1,2$ and $j=a,b$. Although the hospice manager cannot identify the misdiagnosed patients at admission, she can do so during their stay at the hospice. As before, let r and c_i^j denote the daily revenue and cost of caring for a class i,j patient ($i=1,2$, $j=a,b$), respectively. Consider a patient of class i,j admitted at time $t\in[0,T]$, and let $r_i^j(t)$ and $c_i^j(t)$ denote potential revenues to be collected from Medicare and the cost of caring for that patient over the remainder of the fiscal year $[t,T]$, respectively.

At the end of the year, we use a terminal value v_i^j to denote the ongoing value of a class i,j patient ($i=1,2$ and $j=a,b$) who lives into the following fiscal year. Note that this value may be negative if the hospice's revenue is severely constrained by its cap, so that the patient costs money to treat but brings in little or no revenue. Assuming a fixed terminal value for all patients (regardless of when they were admitted) is exact only if LOS distributions are exponential and must otherwise be considered an approximation based on the average LOS. Terminal values are explored further in §4.2. We let $v_i^j(t)$ denote the terminal value attributed to a class i,j patient arriving at time t. The following proposition characterizes $r_i^j(t)$, $c_i^j(t)$, and $v_i^j(t)$.

Proposition 2. *For $t\in[0,T]$, $i=1,2$, and $j=a,b$,*

$$r_i^j(t)=r\int_0^\infty [x\wedge(T-t)]f_i^j(x)\,dx,$$

$$c_i^j(t)=c_i^j\int_0^\infty [x\wedge(T-t)]f_i^j(x)\,dx, \quad \textit{and}$$

$$v_i^j(t)=v_i^j\int_{T-t}^\infty f_i^j(x)\,dx.$$

For ease of notation, we define

$$r_i(t):=\sum_{j=a}^b \gamma_i^j r_i^j(t), \quad c_i(t):=\sum_{j=a}^b \gamma_i^j c_i^j(t),$$

$$v_i(t):=\sum_{j=a}^b \gamma_i^j v_i^j(t).$$

The hospice manager is required to admit all arriving patients but also faces a decision as to whether to actively recruit more patients. Recall that, during admission, the hospice manager cannot identify the patients who may be discharged alive. Let $\alpha_i(t)$ denote the rate at which the hospice manager recruits class i patients at time t. There is a convex increasing cost $s_i(\alpha)$ associated with recruiting (or searching for) class i patients at rate α. For concreteness, assume $s_i(\alpha) = \frac{1}{2}\eta_i^s \alpha^2$ for $i = 1, 2$ and $\alpha \geq 0$, where $\eta_1^s, \eta_2^s > 0$ are given parameters.

Moreover, the hospice manager may choose to live discharge patients who are eligible. We assume for simplicity that the hospice manager learns which patients are eligible for live discharge soon after their admission. Let $\theta_i(t)$ denote the rate at which the hospice manager live discharges class i patients at time t. By live discharging a class i type b patient, the hospice manager forgoes potential revenues of $\tilde{r}_i(t)$ and terminal value of $\tilde{v}_i(t)$ but saves the caring cost of $\tilde{c}_i(t)$ for $i = 1, 2$; these can be calculated as in Proposition 2 given the length-of-stay distributions of class i patients to be live discharged.

The hospice manager is constrained while making the live-discharge decisions in two ways: First, the number of eligible patients may be limiting. Second, discharging additional living patients becomes increasingly more difficult. Our model does not keep track of the evolution of the number of patients in the hospice because doing so would require a complex history-dependent model, which is intractable analytically. This complexity stems from the "memory" of the length-of-stay distributions and prevents us from incorporating the constraint on the number of patients eligible for live discharge.[10] Instead, for tractability, we assume a convex increasing cost $g_i(\theta) = \frac{1}{2}\eta_i^l \theta^2$ associated with the live-discharge rate of θ (for $i = 1, 2$) to capture the difficulties in the live discharges indirectly. Furthermore, Taylor et al. (2008) report that 15.5% of all patients are live discharged, providing evidence that the number of patients eligible for live discharges is not too low. Therefore, replacing hard constraints by a convex increasing cost for live discharges may be a reasonable proxy.

Let $R(0)$ and $C(0)$ denote the potential revenue and the caring cost, respectively, associated with patients in the hospice at the beginning of the fiscal year. Then given the hospice manager's recruiting policy $\alpha(\cdot)$ and the live-discharge policy $\theta(\cdot)$, the cumulative (potential) revenues up to time $t \in [0, T]$, denoted by $R(t)$, are given by

$$R(t) = R(0) + \sum_{i=1}^{2} \int_0^t r_i(s)[\lambda_i + \alpha_i(s)]\, ds - \sum_{i=1}^{2} \int_0^t \tilde{r}_i(s)\theta_i(s)\, ds.$$

Similarly, the cumulative caring cost incurred by the hospice manager up to time t is given by

$$C(t) = C(0) + \sum_{i=1}^{2} \int_0^t c_i(s)[\lambda_i + \alpha_i(s)]\, ds - \sum_{i=1}^{2} \int_0^t \tilde{c}_i(s)\theta_i(s)\, ds,$$

and the cumulative terminal value associated with patients in the hospice at time t is given by[11]

$$V(t) = \sum_{i=1}^{2} \int_0^t v_i(s)[\lambda_i + \alpha_i(s)]\, ds - \sum_{i=1}^{2} \int_0^t \tilde{v}_i(s)\theta_i(s)\, ds.$$

The cumulative recruiting costs $S(t)$ and the live-discharge cost $D(t)$ up to time t are given by

$$S(t) = \sum_{i=1}^{2} \int_0^t s_i(\alpha_i(x))\, dx \quad \text{and} \quad D(t) = \sum_{i=1}^{2} \int_0^t g_i(\theta_i(x))\, dx. \tag{2}$$

As mentioned earlier, a crucial feature of the Medicare reimbursement policy is that the hospice's revenue is constrained by a cap, which increases with the number of patients admitted during the fiscal year. To be specific, the cap at time t is given by

$$K(t) = K \sum_{i=1}^{2} \int_0^t (\lambda_i + \alpha_i(x))\, dx. \tag{3}$$

Therefore, the realized revenue at the end of the fiscal year is given by $\min\{K(T), R(T)\}$, and the hospice manager's problem (P) can be written as follows: Choose search rates $\alpha_i(\cdot)$ and live-discharge rates $\theta_i(\cdot)$ for $i = 1, 2$ dynamically so as to

$$\begin{aligned} &\text{maximize} \;\; \min(K(T), R(T)) + V(T) - C(T) \\ &\qquad\qquad - S(T) - D(T) \\ &\text{subject to} \;\; \alpha_i(t) \geq 0, \theta_i(t) \geq 0 \quad \text{for all } i, t. \end{aligned} \tag{P}$$

Although the hospice manager's problem (P) is an optimal control problem (and hence, infinite dimensional), its dual is much simpler. In Online Appendix B, we show that the dual formulation can

[10] The Center for Medicare and Medicaid Studies (2010) states that any remaining portion of the cap the discharged living patient contributed to the hospice can remain at the hospice unless that patient elects to transfer to another hospice immediately or later (the originating hospice would then relinquish that part of the cap to the other hospice). Because the patients being discharged alive in our model are leaving because of "improved or stabilized" conditions (Kutner et al. 2004), this would not apply to our model. This discharge reason is the only one under the responsibility of the provider, whereas the other reasons (e.g., patient/family decision, pursuit of more aggressive treatment, or transfer to another hospice; see Kutner et al. 2004) are those of the patient and are likely to result in a cap reduction.

[11] We assume that those patients who are in the hospice at the beginning of the fiscal year are not present at the end as a result of death or being discharged alive. Therefore, those patients do not contribute to the terminal value $V(T)$.

be reduced to a one-dimensional convex optimization problem, enabling an explicit solution to both the dual and the hospice manager's original problem. The proof of this relies on the duality theory for optimal control problems developed by Rockafellar (1970), which is also introduced in Online Appendix B. For $q \in [0, 1]$, define

$$F(q) = -K(\lambda_1 + \lambda_2)T + R(0) + \sum_{i=1}^{2} \lambda_i \int_0^T r_i(s)\,ds + \sum_{i=1}^{2} \int_0^T \frac{-\tilde{r}_i(t)}{\eta_i^l} [-\tilde{r}_i(t)q + \tilde{c}_i(t) - \tilde{v}_i(t)]^+\,dt + \sum_{i=1}^{2} \int_0^T \frac{r_i(t) - K}{\eta_i^s} \cdot [K + (r_i(t) - K)q - c_i(t) + v_i(t)]^+\,dt. \quad (4)$$

We now assume that the following holds:

$$\min\{r_i(t), K\} + v_i(t) - c_i(t) > 0 \quad \text{for some } i,\, t \in (0, T), \quad (5)$$

which ensures the strict monotonicity of F and the uniqueness of the dual optimal solution.

In Online Appendix B, $F(\cdot)$ is shown to be the derivative of the dual objective function. The following proposition shows that the inverse F^{-1} of F is well defined.

Proposition 3. *F is continuously differentiable and strictly increasing.*

We are now ready to state our main result.

Theorem 1. *If $F(0) > 0$, then let $q^* = 0$; if $F(1) < 0$, then let $q^* = 1$. Otherwise, let $q^* = F^{-1}(0)$. Then the hospice manager's optimal recruiting and live-discharge rates for $i = 1, 2$ and $t \in [0, T]$ are given by*

$$\alpha_i^*(t) = \frac{[K(1 - q^*) + q^* r_i(t) - c_i(t) + v_i(t)]^+}{\eta_i^s}$$

and

$$\theta_i^*(t) = \frac{[-\tilde{r}_i(t)q^* + \tilde{c}_i(t) - \tilde{v}_i(t)]^+}{\eta_i^l},$$

respectively.

We have thus explicitly characterized the optimal recruiting and live-discharge rates for a given hospice. For either type of patient, the hospice may choose to never actively recruit those patients, relying entirely on their natural arrival rates; they may recruit patients throughout the year, only at the beginning of the year; or they may only recruit toward the end of the Medicare year. Similar patterns apply to live discharging. We now exercise these findings numerically.

4.2. Numerical Study of the Fluid Model

In the numerical experiments in this subsection, the life spans of patients will be modeled as gamma distributions. A randomized sample of the 1993 cohort of Medicare hospice beneficiaries (184,843 data points) was obtained, including the number of days between the date of admission and the date of death for 27 disease categories. The histograms describing the number of days survived since admission for each disease were monotone decreasing (in time "buckets" of one day), each characterized by very high frequencies at low numbers of days and decreasing into what could be described as "light tails." Each individual disease's histogram was curve-fitted, and we also grouped diseases into type 1 (all cancer diseases and end-stage renal failure) and type 2 (all noncancer diseases other than renal failure) and fitted those histograms. There are a number of parametric distributions that tended to fit these histograms quite well, and the gamma distributions tended to be consistently good performers for many diseases and across various measures of fit. The gamma distribution has the added benefit that the mixing of such distributions (for a common scale parameter) will also be a gamma distribution, which serves the purpose of separating those distributions for subtypes a and b, resulting in $\Gamma_1^a(0.1567, 455.14)$, $\Gamma_1^b(0.8777, 455.14)$ and $\Gamma_2^a(0.2198, 550.61)$, $\Gamma_2^b(1.2308, 550.61)$ (these are the shape and scale parameters for types 1 and 2, respectively).

For the purposes of our examples, we doubled the scale parameter (to 1,101.22) for the type 2 distributions to represent those hospices experiencing longer distributions for type 2 patients, where binding cap scenarios are more likely. There are a couple of reasons for why the fitted distributions may not result in binding cap scenarios. First, these data are for a randomized sample of patients drawn from the entire country and will not reflect a binding cap scenario for a single hospice (the cap is *designed* not to bind for an average hospice). Second, these data are from the 1993 cohort, which was before the upper tails of the type 2 distributions began elongating, so such distributions are unlikely to contribute to a binding cap.

The study by Taylor et al. (2008) contains data where the proportion of patients who enter hospice and leave before death is 11.3%.[12] In the absence of any other data, we use these data to set $\gamma_1^b = 1 - \gamma_1^a = \gamma_2^b = 1 - \gamma_2^a = 0.113$, reflecting the proportion of incorrectly diagnosed (or *not* terminally ill) patients

[12] Taylor et al. (2008) report that among the 1,218 patients under study, 1,029 died under continuous care, 131 were discharged alive and not readmitted, and 58 were live discharged and then readmitted and died under hospice care (a few patients in this third group were discharged a second time but then died outside of hospice care). We excluded this third category because we wished to only examine those patients who were discharged alive a single time and then died. Thus, our proportion of potentially live-dischargable patients is $131/(131 + 1{,}029) = 11.3\%$.

Figure 1 Optimal Recruiting and Live-Discharge Rates for Various Values of $R(0)$ for $\psi = 0.15$

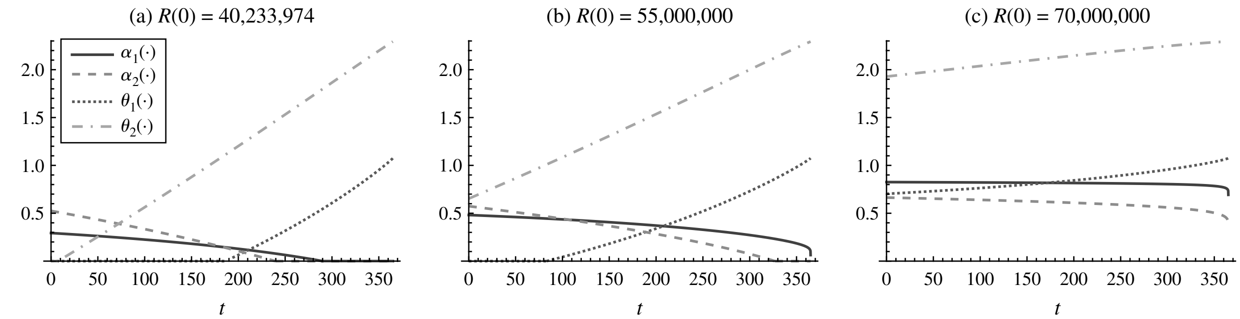

at admission. The Medicare parameters for the 2007 Medicare year are $K = \$21{,}410$ and $r = \$130.79$. The natural daily arrival rates of patient are chosen as $\lambda_1 = 1$ patient per day and $\lambda_2 = 9$ patients per day, intended to reflect the proportions of a hospice experiencing cap issues.[13] We have taken the daily cost of caring as $c_1 = \$60$ per day and $c_2 = \$45$ per day from Killaly and Mukamel (2010). The recruiting and live-discharging parameters are $\eta_i^j = 15{,}000$ for $i = 1, 2$ and $j = l, s$. Varying these parameters results in greater or lesser recruiting and live-discharging rates, but no additional insight. We set the terminal values as $v_i^j = (\psi r - c_i)m_i^j$, where the exogenous parameter ψ is exercised between 0 and 1, as a robustness check. Note that, for all but type 2b, the average life of patients living into the next year is actually less than m_i^j because the given gamma distribution has an increasing failure rate, so $-c_i m_i^j$ is a lower bound on the terminal value and $(r - c_i)m_i^j$ is an upper bound on the terminal value. For ease of calculation, we use $\tilde{r}_i(t) = r_i^b(t)$, $\tilde{c}_i(t) = c_i^b(t)$, and $\tilde{v}_i(t) = v_i^b(t)$.

In the following examples, we examine the optimal recruiting and live-discharge policies of the hospice manager for a variety of $R(0)$ and ψ values. Varying ψ is necessary because it does not seem possible to determine precisely the ongoing value of those patients who are alive at the end of the Medicare year. We have, however, determined a base-case value of $\psi = 0.15$ by using a steady-state simulation as described in §4.3. Parameter $R(0)$ is also varied to reflect that a hospice will begin a new year with living patients who represent a stream of revenues in the new year. The hospice will adjust its optimal recruiting and live discharging for varying values of $R(0)$ and ψ. Of course, $R(0)$, $C(0)$, and ψ will be related; they reflect the amount of revenue and costs patients enrolled at the start of the year accrue, and the value of a patient at the end of the year, respectively. Finding $\psi = 0.15$ is our attempt to relate $R(0)$ and ψ, but this approach is imperfect, so we exercise our model numerically across these parameters.

In Figure 1 we see the effect of varying $R(0)$[14] for $\psi = 0.15$. At lower values of $R(0)$ (see Figure 1(a)), the cap is not binding, and we observe the hospice will recruit type 2 patients at a higher rate than type 1 patients (the daily margin of patient type 1 is $130.79 - 60 = \$70.79$ and of patient type 2 is $130.79 - 45 = \$85.79$), although both patient types will be sought, which is not surprising as each patient type is profitable. As the cap begins to bind (see Figure 1(b)), we see that the hospice extends the duration of its recruiting, and the recruiting of type 1 patients is sustained longer at a higher level and begins to dominate that of the type 2 patients during the latter part of the year. This stems from the fact that the patients at the end of the year are costly because $\psi = 0.15$ ($v_i < 0$ for $i = 1, 2$), and type 1 patients can bump up the cap as much as type 2 patients but live shorter durations and thus will be less costly if they happen to remain living at the end of the year. (Clearly, many of them do not survive until then unless they are recruited very close to the end of the fiscal year.) This effect continues as $R(0)$ increases, as we see in Figure 1(c). Note, however, that the fact that all patients who live into the following year are costly (regardless of their number) is an artifact of our specific model and choice of $\psi = 0.15$.

We also observe that the hospice will discharge living patients at a greater rate toward the end of the year to reduce the chance of having these costly patients by year's end. Moreover, by live discharging patients, the hospice manager also avoids the caring costs for those patients during the current fiscal year. Because the cost of live discharging is assumed to be convex increasing, the hospice will conduct this

[13] Note that total arrival rates of 10 patients per day would reflect a relatively large hospice (Wright and Katz 2007), but the simulation model that will be presented shortly required arrival rates of this magnitude to avoid issues with integrality. However, the fluid model and intuition given later is robust to the choice of λ.

[14] The values of $R(0)$ in Figure 1 can be mapped to a number of living patients enrolled in the hospice at the beginning of the year. For example, \$40.2 million translates to approximately 2,171 patients, where this number was estimated using a steady-state simulation as described in §4.3. Note that natural arrivals of patients are $(\lambda_1 + \lambda_2)365 = 3{,}650$ patients per year, and Little's law applied to a stationary hospice would imply 3,419 patients.

Figure 2 **Optimal Recruiting and Live-Discharge Rates for Various Values of ψ for $R(0) = 55{,}000{,}000$**

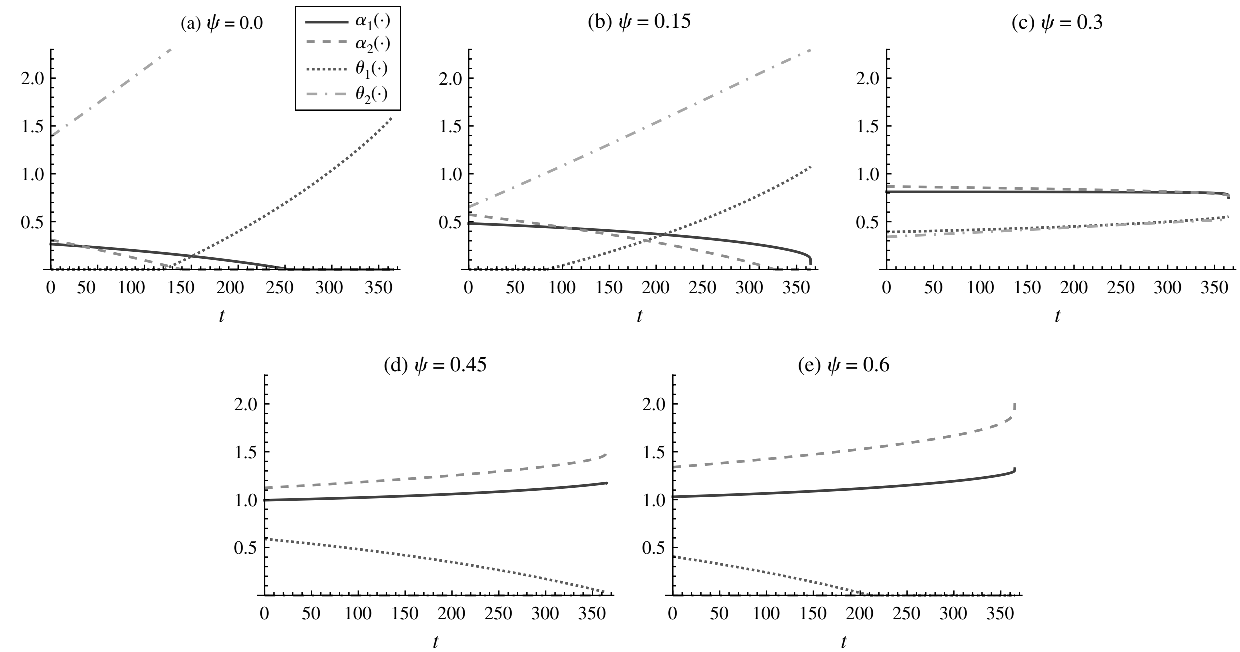

over a period of time toward the end of the year rather than all at once on the final day, for example. As $R(0)$ continues to increase, these trends will continue: live discharges will continue over a longer period, with type 2 discharges dominating type 1 later in the year (because $v_2 < v_1 < 0$), and discharge rates will increase. This is a means of reducing these initial patients who contribute nothing positive toward the cap but who would consume it if they remained under care. As one would expect, for much lower values of $R(0)$ when the cap is not binding, little discharging alive occurs. This is because the hospice has little need to reduce its patient census and can accommodate the revenues from those patients under the cap (to which they do not contribute at all). For $\psi = 0.15$ and $R(0) = 55{,}000{,}000$, we find that 9.1% of all patients are discharged alive, which is comparable to the numbers found by Kutner et al. (2004) and others. Clearly, $R(0)$, which represents the potential revenues of patients left over from the previous year, has a considerable effect on the optimal actions of the hospice. Compensating actions such as additional recruiting (to increase the cap) or additional discharging of living patients (to reduce the revenues) are clearly necessary as $R(0)$ increases.

In Figure 2 we observe the effect of increasing ψ for a fixed value of $R(0)$.[15] When $\psi = 0$, $v_1 = -6{,}503.28$ and $v_2 = -16{,}548.9$, so the recruiting of type 1 patients is higher than that of type 2, and both will decrease toward the end of the year. Likewise, the terminal costs are so significant that sizable live-discharging will take place, with more type 2 patients discharged than type 1. As ψ increases, the gap between v_1 and v_2 narrows. As this happens, the rate of type 2 patient recruiting will increase, eventually dominating the rate of type 1 patient recruiting, both eventually increasing at the end of the year. Likewise, as ψ increases and the terminal values become more positive (that is, there is no negative effect of carrying patients into the following year), the need for discharging alive reduces significantly and ultimately disappears.

The manager's objective is to ensure that the hospice is not exposed to a large deficit in the difference between the cap and revenues received. This is verified by examining Figure 3, which is for $R(0) =$ \$55 million and $\psi = 0.15$. In Figure 3(a) the difference between the cap and the revenues for the natural arrivals (i.e., absent any recruiting) is displayed, demonstrating a shortfall of \$13.8 million at time T. Once the hospice's recruit and live-discharge actions are included, the gap between the cap and the revenues, displayed in Figure 3(b), will be closed at time T. The last result we wish to highlight is that when we increase the coefficient of variation of the LOS distribution, the recruiting and live-discharge rates for both disease types increases.

Thus far in this section, we have assumed a deterministic model. A natural question to ask is, what are the effects of random arrivals and LOS on the

[15] The terminal parameters are as follows: (a) $v_1 = -6{,}503.3, v_2 = -16{,}548.9$; (b) $v_1 = -4{,}376.9, v_2 = -9{,}334.2$; (c) $v_1 = -2{,}250.5$, $v_2 = -2{,}119.4$; (d) $v_1 = -124$, $v_2 = 5{,}095.4$; and (e) $v_1 = 2{,}002.4$, $v_2 = 12{,}310.2$. The line for $\theta_2(t)$ in Figure 2(a) continues to increase and is truncated so that the scale is the same on all plots.

Figure 3 The Gap Between the Cap and Potential Revenues ($R(0) = \$55$ **million**, $\psi = 0.15$)

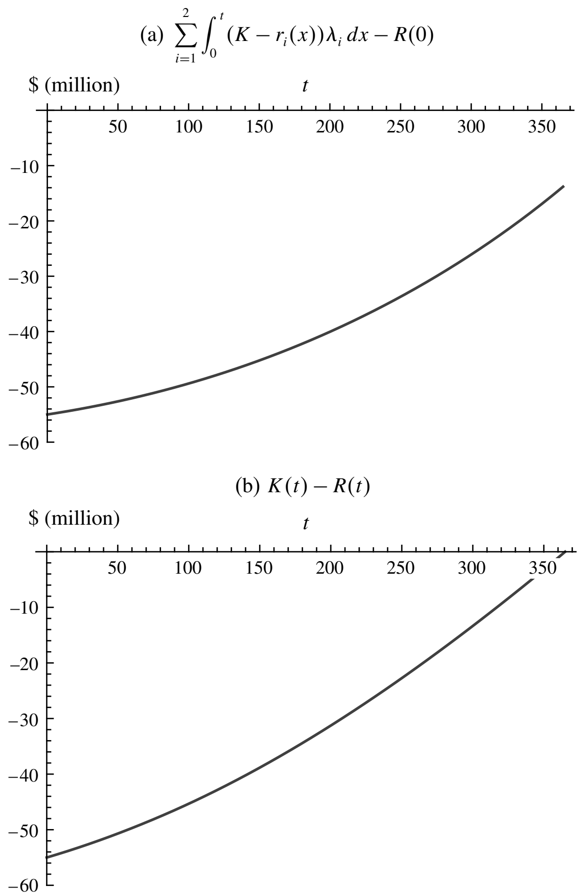

intuition provided? The following section provides a simulation study to address this question.

4.3. A Simulation Study

To test the effects of stochasticity on the system, we created a discrete event simulation in C++. Although using this simulation to find optimal recruiting and discharge policies is not tractable, we can use it to test some reasonable hypotheses for the system. In particular, we hypothesize that uncertainty should amplify the effects of the cap as a result of random shocks, resulting in more type 1 patients being recruited, especially toward the end of the horizon, and perhaps decreasing type 2 patient recruiting.

The simulation was created to mimic the fluid model except with the addition of stochastic arrivals and LOS. In particular, multiple replications of a single year (i.e., terminating) model were performed, where the terminal values were set equal to those of the numerical study of the fluid model and an initial number of customers were introduced at the start of each year to generate $R(0)$ and $C(0)$. However, because terminal values and initial customers are actually endogenous to the system, in addition to this "terminating simulation," we also created a steady-state version. Admissions were assumed to take place at the start of each day, with a random number of patients arriving at a rate equal to the sum of the natural arrivals and the recruiting rate. To test the effects of variability, the number of arrivals each day was sampled from a discretized gamma distribution with coefficient of variation set to 1 for the base case (in which case the number of arrivals each day follows a discretized exponential distribution). LOS distributions were also taken from a gamma distribution with parameters as given previously for the fluid model.

To test the above hypotheses, a heuristic was created based on state-dependent control (recall that finding the optimal control in the stochastic setting is not tractable). For each day, the hospice manager wishes to decide recruiting and discharge rates in order to maximize expected revenue minus costs over what is left of the horizon. As in the fluid model, this is implemented by solving an optimization problem on the revenue, cost, and terminal value functions, where the decision variables are the recruiting and discharge rates for the remaining days in the year. Only the calculated recruiting and discharge rates for the current day are used. Those for the rest of the year are discarded (and recalculated the following day). Further details on the heuristic may be found in Online Appendix C. We ran the heuristic on all parameter choices in Figures 1 and 2, and the average improvement over the fluid control was 6.29%, illustrating that there is typically a benefit to state-dependent control in the presence of stochasticity. As mentioned previously, we also created a steady-state simulation using the heuristic that imputes terminal values.[16]

In Figure 4 we observe the difference between the terminating-simulation heuristic recruiting and the recruiting in the fluid policy, and the difference between the terminating-simulation heuristic discharging and the discharging in the fluid policy for $\psi = 0.15$ and 2,968 initial patients split appropriately between the different classes, which corresponds to an $R(0)$ of approximately \$55 million.[17] Note that the rates shown in Figure 4 are averages over all

[16] We ran the steady-state simulation with the heuristic policy under the base-case parameters from the fluid numerics section to create a base-case estimate for both the terminal values and the number of customers present at the start of the year. At the end of each year, the estimates for the terminal values are updated using an algorithm described in Online Appendix C. Note that terminal values are used within the heuristic optimization to determine the appropriate levels of recruiting and/or discharging. Therefore an iterative approach is taken to their computation.

[17] We also produced these graphs for 2,171 initial patients, which are the values obtained from the steady-state simulation and correspond to an $R(0)$ of approximately \$40.2 million, but we found in that case the heuristic did almost no discharging and the figure was less informative.

Figure 4 Difference in Recruiting and Live-Discharge Rates Between the Heuristic and the Simulated Fluid Policy for $R(0) = 55{,}000{,}000$ **and** $\psi = 0.15$

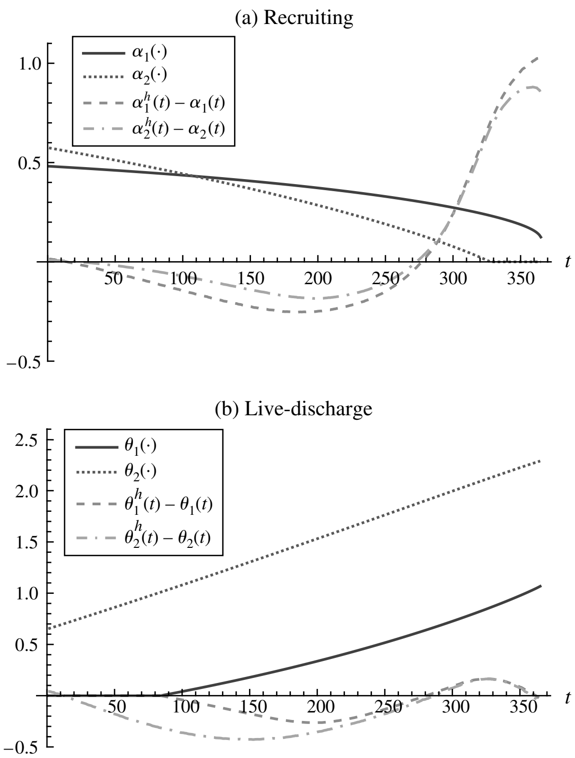

Note. Superscript *h* is the heuristic.

replications. In general, the confidence intervals for these rates increased in width throughout the year; this is intuitive because the effects of stochasticity would tend to accumulate throughout the year and in some years no recruiting would be desirable by the end, whereas in others, much recruiting may be needed.

Figure 4(a) shows that in the simulation we see more recruiting of both types of patients toward the end of the horizon. This supports our hypothesis that in a stochastic system, cancer patients may be more heavily sought, particularly at the end of the year when revenues may have exceeded the cap. However, there is also more recruiting of noncancer patients toward the end of the year, contradicting our initial hypothesis that there should be less recruiting of these patients. On the other hand, the difference between the heuristic and fluid recruiting of type 1 patients is greater than that of type 2 patients at the end of the year, again supporting our intuition.

In Figure 4(b), we see that discharging living patients under the heuristic is at first less and then more than the fluid model, with a lower overall average level. One possible reason is that the stochastic model may delay discharging until uncertainty has been further resolved. As noted earlier, with fewer patients at the start of the year ($R(0) = \$40.2$ million), the heuristic does far less discharging than the fluid model, although the decreased discharging of patients was paired with increased recruiting by the heuristic. Given the possibility of “favorable” sample paths under stochasticity, where no discharging is needed, this appears to be consistent with the delayed discharging decisions shown in Figure 4(b).

To further test the effects of stochasticity, we ran the terminating simulation with $\psi = 0.15$ under a variety of values for the squared coefficient of variability of both the number of arrivals per day and the LOS distributions. Specifically, we scaled all squared coefficients of variability up by 10 times and down by 10 times while keeping the means constant. We found that the improvement over the fluid control policy increased with increased variability, verifying our intuition that a state-dependent policy, such as the heuristic, will have more value as the system gets more variable. Furthermore, there was a clear trend of profit decreasing and both recruiting and discharging levels of both types of patients increasing with increasing variability of arrivals. When the squared coefficients of variability of the LOS distributions were altered, both recruiting and discharge levels increased significantly under increased variability. For this case, similar trends can also be generated using the fluid model because it allows for two-parameter LOS distributions.

We believe that nonstationary recruiting and live-discharge rates are undesirable from a policy perspective because they do not align with the equal access objective of a social planner. The immediate implication of nonstationary recruiting is that a prospective patient’s desirability to a hospice depends on the calendar. A cancer patient would not be as highly sought in December as he or she might be in October (the Medicare year runs from November 1 to October 31). Similarly, the nonstationary discharge of living patients is highly undesirable. The government’s motivation in establishing the Medicare hospice benefit was to facilitate comfortable palliative care for patients during their end-of-life by third-party providers. Nonstationary active recruiting and live discharges, which can arise as the optimal policy and are observed in practice, are clearly unintended consequences of the government’s Medicare reimbursement policy. The following section presents a new policy for Medicare to consider. This policy is intended to alleviate this unintended behavior.

5. The Legacy Policy and Its Analysis

This section introduces and analyzes a policy to overcome the nonstationary behavior observed under the current Medicare reimbursement policy seen in §4. We label this the *legacy policy* because it explicitly addresses the possibility of beneficiaries living into the next Medicare year. The implementation of this policy requires the hospice to segregate the tracking of the patients admitted in each Medicare year until they expire, but it is otherwise no more burdensome in terms of administration than the current Medicare policy. Indeed, our goal is to create a policy that maintains the same fundamental framework of the Medicare policy but does not have an inherent incentive for nonstationary recruiting and discharge. That is, we assume that Medicare wishes to maintain a cap (to limit its exposure, especially as the declaration of a patient being terminally ill can be difficult), to keep the cap pooled (to mitigate risk to the hospice), and to keep payment rates constant across disease types and time frame (for ease of implementation).

The legacy policy consists of allowing the hospice to continue receiving revenues for all the patients living at the end of the year until any remaining cap is exhausted or all these patients expire, whichever occurs first. If the cap is exhausted before these patients expire, then the hospice receives no further revenues for these patients but continues to incur the costs of caring for them. So the key difference between the legacy policy and the current Medicare policy is that once a new Medicare year arrives, current patients (those receiving hospice services currently) will count against the previously accumulated cap (if it is not exhausted) as opposed to the current practice of counting them against the cap accumulated in the new year. In practice, this policy is likely most simply implemented by allowing exactly one extra legacy year and assuming that the probability of having patients live into a third Medicare year with the cap still not exhausted is negligible.[18]

Under the modified policy, because the accounting is over the entire patient life rather than just over the Medicare year, $r_i(t) = rm_i$, $c_i(t) = c_i m_i$ and $\tilde{r}_i(t) = r\tilde{m}_i$, $\tilde{c}_i(t) = c_i\tilde{m}_i$ for all i, t, where $\tilde{m}_i$ denotes the mean of the residual life span of patients who are discharged alive. Similarly, $R(0) = C(0) = 0$, $v_i = \tilde{v}_i = 0$, and $v_i(t) = \tilde{v}_i(t) = 0$ for all i, t. Then the fluid model of the hospice manager's problem can be adapted to choosing recruiting rates $\dot{z}(\cdot)$ and live-discharge rates $\dot{\Theta}(\cdot)$ to

$$\max\min\left\{K\sum_{i=1}^{2}(\lambda_i T + z_i(T)),\; r\sum_{i=1}^{2}(m_i(\lambda_i T + z_i(T)) - \tilde{m}_i\Theta_i(T))\right\} - \sum_{i=1}^{2}[c_i m_i(\lambda_i T + z_i(T)) - c_i\tilde{m}_i\Theta_i(T)]$$

$$-\sum_{i=1}^{2}\int_0^T s_i(\dot{z}_i(t))\,dt - \sum_{i=1}^{2}\int_0^T g_i(\dot{\Theta}_i(t))\,dt \tag{6}$$

subject to

$$z_i(t) = \int_0^t \dot{z}_i(s)\,ds \quad \text{with } \dot{z}_i(t) \geq 0 \text{ for all } i, t, \tag{7}$$

$$\Theta_i(t) = \int_0^t \dot{\Theta}_i(s)\,ds \quad \text{with } \dot{\Theta}_i(t) \geq 0 \text{ for all } i, t. \tag{8}$$

The following result significantly simplifies the above control problem.

Proposition 4. *Any optimal solution of the formulation* (6)–(8) *is a stationary recruiting policy; i.e.,* $\alpha_i(t) = \dot{z}_i(t) = \bar{\alpha}_i$ *and* $\theta_i(t) = \dot{\Theta}_i(t) = \bar{\theta}_i$, *where* $\bar{\alpha}_i \geq 0$, $\bar{\theta}_i \geq 0$ *for all* i, t.

Therefore, the hospice's manager's problem can be stated without loss of generality (setting $T = 1$) as follows: Choose the stationary recruiting rates α_1, α_2 and the stationary live-discharge rates θ_1, θ_2 so as to

$$\max\left\{\min\left[K\sum_{i=1}^{2}(\lambda_i + \alpha_i),\, r\sum_{i=1}^{2}[(\lambda_i + \alpha_i)m_i - \tilde{m}_i\theta_i]\right] - \sum_{i=1}^{2}c_i[m_i(\lambda_i + \alpha_i) - \tilde{m}_i\theta_i] - \sum_{i=1}^{2}s_i(\alpha_i) - \sum_{i=1}^{2}g_i(\theta_i)\right\} \tag{9}$$

subject to $\alpha_i \geq 0$, $\theta_i \geq 0$, $i = 1, 2$.

To facilitate the statement of this section's main result, define

$$\pi_1 = \left(\sum_{i=1}^{2}\frac{(rm_i - K)^2}{\eta_i^s} + \sum_{i=1}^{2}\frac{r(\tilde{m}_i)^2 c_i}{\eta_i^l} - \sum_{i=1}^{2}(rm_i - K)\left[\lambda_i + \frac{(r - c_i)m_i}{\eta_i^s}\right]\right) \cdot \left(\sum_{i=1}^{2}\frac{(rm_i - K)^2}{\eta_i^s} + \sum_{i=1}^{2}\frac{(r\tilde{m}_i)^2}{\eta_i^l}\right)^{-1},$$

$$\pi_2 = \left(\sum_{i=1}^{2}\frac{(rm_i - K)^2}{\eta_i^s} + \frac{r(\tilde{m}_1)^2 c_1}{\eta_1^l} - \sum_{i=1}^{2}(rm_i - K)\left[\lambda_i + \frac{(r - c_i)m_i}{\eta_i^s}\right]\right) \cdot \left(\sum_{i=1}^{2}\frac{(rm_i - K)^2}{\eta_i^s} + \frac{(r\tilde{m}_1)^2}{\eta_1^l}\right)^{-1},$$

$$\pi_3 = \left(\sum_{i=1}^{2}\frac{(rm_i - K)^2}{\eta_i^s} - \sum_{i=1}^{2}(rm_i - K)\left[\lambda_i + \frac{(r - c_i)m_i}{\eta_i^s}\right]\right) \cdot \left(\sum_{i=1}^{2}\frac{(rm_i - K)^2}{\eta_i^s}\right)^{-1}.$$

The following lemma will be useful in proving our main result.

[18] Note that this policy would mean that a start-up hospice would have no cap adjustment to pay back until the end of its second year of operation.

LEMMA 1. *Assume $c_1 > c_2$. Then the following hold*:
(i) $\pi_1 < c_2/r$ *if and only if* $\pi_2 < c_2/r$.
(ii) *If* $\pi_1 < c_2/r$, *then* $\pi_3 < c_1/r$.
(iii) $\pi_2 < c_1/r$ *if and only if* $\pi_3 < c_1/r$.

We are now ready to state the main result of this section.

PROPOSITION 5. *Let* α_i^*, θ_i^* *for* $i = 1, 2$ *be the optimal solution to* (9). *Then*

(1) *If* $\sum_{i=1}^2 (rm_i - K)[(r-c_i)m_i/\eta_i^s + \lambda_i] < 0$, *then*

$$\alpha_i^* = \frac{m_i}{\eta_i^s}(r - c_i) \quad \text{and} \quad \theta_i^* = 0 \quad \text{for } i = 1, 2.$$

(2) *If*

$$\sum_{i=1}^2 (rm_i - K)[(r-c_i)m_i/\eta_i^s + \lambda_i] > \sum_{i=1}^2 (rm_i - K)^2/\eta_i^s + \sum_{i=1}^2 r(\tilde{m}_i)^2 c_i/\eta_i^l,$$

then

$$\alpha_i^* = \frac{K - c_i m_i}{\eta_i^s} \quad \text{and} \quad \theta_i^* = \frac{c_i \tilde{m}_i}{\eta_i^l} \quad \text{for } i = 1, 2.$$

(3) *If* $0 \le \sum_{i=1}^2 (rm_i - K)[(r-c_i)m_i/\eta_i^s + \lambda_i] \le \sum_{i=1}^2 (K - rm_i)^2/\eta_i^s + \sum_{i=1}^2 r(\tilde{m}_i)^2 c_i/\eta_i^l$, *then we have the following cases*:

(a) *If* $\pi_1 < c_2/r$, *then*

$$\alpha_i^* = \frac{K(1-\pi_1) + rm_i\pi_1 - c_i m_i}{\eta_i^s} \quad \text{and} \quad \theta_i^* = \frac{\tilde{m}_i(c_i - r\pi_1)}{\eta_i^l} \quad \text{for } i = 1, 2.$$

(b) *If* $\pi_1 \ge c_2/r$, *then we have the following two subcases*:

(i) *If* $\pi_2 < c_1/r$, *then* $\theta_1^* = m_1(c_1 - r\pi_2)/\eta_1^l$, $\theta_2^* = 0$ *and*

$$\alpha_i^* = \frac{K(1-\pi_2) + rm_i\pi_2 - c_i m_i}{\eta_i^s} \quad \text{for } i = 1, 2.$$

(ii) *Otherwise, i.e.,* $\pi_2 \ge c_1/r$, *then*

$$\alpha_i^* = \frac{K(1-\pi_3) + rm_i\pi_3 - c_i m_i}{\eta_i^s} \quad \text{and} \quad \theta_i^* = 0 \quad \text{for } i = 1, 2.$$

Depending on the parameters, Proposition 5 has three cases for the solutions of the optimal recruiting rate under the legacy policy.

- In case 1, the reimbursements determine the revenue rate (i.e., the cap does not bind).
- In case 2, the cap binds in the optimal solution (i.e., the cap determines the revenue rate).
- In case 3, the cap equals the reimbursement rate (perfectly balanced).

In case 1 of Proposition 5, the potential revenues of the hospice will not use the available cap, and the recommended recruiting rate will then simply solve the resulting first-order condition of (9). This results in rates that balance out the contribution over the remaining life of the patient against the cost of recruiting them, and there is no need to discharge anyone living. These rates are increasing in r and m_i and decreasing in η_i and c_i. Note that if the cap also does not bind under the original policy, then we would expect to see similar levels of recruiting (and zero discharges). Thus, the legacy policy is designed for hospices where the cap is relevant to the revenues and will make little difference to hospices that typically do not meet their caps.

In case 2 the cap is binding, and the resulting optimal recruiting rate takes the cap as the total revenue, subtracts the total costs of caring, and balances this quantity against the cost of recruiting these patients. The rate of discharging living patients balances the remaining lifetime costs of those patients against the cost of discharging them. Thus, the recruiting rates are increasing in K but decreasing in η_i, m_i, and c_i. Note that depending on the cost structure, the live-discharge rate can be quite high.

In case 3 the recruiting and live-discharge rates of cases 1 and 2 are blended in such a way as to create a perfectly balanced situation where the weighted total revenues equals the available cap.

Note that although Proposition 5 states that the legacy policy will deliver a stationary policy, it does not imply that the recruiting rates will be identical across diseases or even equal in proportion to volumes. Indeed, Proposition 5 confirms that the recruiting rates will differ according to the underlying characteristics and economics for each disease. So although the legacy policy addresses the problem of nonstationarity of the existing reimbursement policy, it does not address the issue of the relative profitability across diseases (accounting for the cost of care, the cost of recruiting, and the life spans of patients), which causes different rates of recruiting and live discharges between diseases. This is consistent with the existing policy and could only be addressed by Medicare instigating a disease-specific reimbursement policy. Indeed, Killaly and Mukamel (2010) recommend revisiting Medicare's disease-invariant per-diem reimbursement, but we believe a disease-dependent payment rate would need to be very carefully chosen. Otherwise, the system will be prone to gaming, such as hospices classifying admissions based on the higher margin disease in the case of comorbid patients. Thus, more careful consideration is merited.

Table 1 presents the legacy policy's optimal recruiting (α_i^*) and live-discharge (θ_i^*) rates for $\psi = 0.15$

Table 1 The Legacy Policy Optimal Stationary Recruiting (α_i^*) and Live-Discharge (θ_i^*) Rates, and Time Average of the Optimal Recruiting and Live-Discharge Rates for Various Levels of the Recruiting and Live-Discharge Costs, η_i^j for $i = 1, 2$ and $j = s, l$ ($\psi = 0.15$, $R(0) = 55{,}000{,}000$)

η_i^j	α_1^*	α_2^*	θ_1^*	θ_2^*	$(1/T)\int_0^T \alpha_1(t)\,dt$	$(1/T)\int_0^T \alpha_2(t)\,dt$	$(1/T)\int_0^T \theta_1(t)\,dt$	$(1/T)\int_0^T \theta_2(t)\,dt$
1,000	12.46	13.89	6.298	1.04	2.065	2.782	3.558	16.53
5,000	2.530	2.635	1.539	1.154	0.4131	0.5563	0.7115	3.306
10,000	1.289	1.228	0.9436	1.168	0.2917	0.3157	0.3926	1.808
15,000	0.8757	0.7596	0.7452	1.173	0.3649	0.2934	0.3544	1.463
20,000	0.6689	0.5252	0.6461	1.176	0.3813	0.2906	0.3724	1.256

and $R(0) = \$55$ million and various recruiting (η_i^s) and live-discharge (η_i^l) costs (for $i = 1, 2$). For comparison purposes, the time averages of the fluid model's optimal recruiting and live-discharge rates are also included (the final four columns). There are several observations we can make based on the table. As might be expected, as the costs of recruiting increase, the legacy policy's recruiting rates decrease because it becomes more expensive to seek patients. However, as the costs of live discharging increase, the live discharge of cancer patients falls while the live discharge of noncancer patients rises. Note that all these costs are increasing concurrently, suggesting that there is a degree of substitution between the live discharges. Moreover, for low values of η_i^s, the legacy policy's recruiting rates start as $\alpha_1^* < \alpha_2^*$, but as η_i^s increase, the fluid model migrates to a cap-constrained scenario. We likewise see that the legacy policy migrates to $\alpha_1^* > \alpha_2^*$ (in Table 1's example, we see that the gap between the fluid model's average recruiting rates is initially large and then narrows). The reverse happens with the legacy policy's live-discharge rates: at low values of η_i^s, $\theta_1^* > \theta_2^*$, but this reversed to $\theta_1^* < \theta_2^*$ at higher levels. This suggests that another substitution effect is going on: the legacy policy is substituting recruiting for the live discharge of cancer patients as the costs increase but the reverse (the live discharge for recruiting) of noncancer patients.

Note also that the legacy policy recommends higher recruiting and type 1 live-discharge rates and lower type 2 live-discharge rates than the average numbers from the fluid model for these negative terminal values. The first suggestion is that although the legacy policy overcomes the undesirable nonstationarity of the fluid model, it appears to encourage other behaviors that Medicare may find undesirable, although this insight is tempered by the knowledge that the fluid model averages may be somewhat depressed by their negative terminal values, quantities that the legacy policy is not subject to. Both recruiting and live discharge are actions that Medicare finds concerning.

Note that the revenues accrued under the legacy policy are greater than those of the fluid policy, and hence the cost to Medicare has increased under this policy. For example, for $\psi = 0.15$, $R(0) = \$55$ million, and $\eta_i^j = 15{,}000$, the legacy policy accrues revenues of \$90,926,088, a 9.2% increase over the fluid policy's revenues of \$83,290,870. Again, this must be tempered by the observation that the comparison is not completely analogous because of the terminal values of the fluid model. Obtaining an accurate estimate of both the revenue effects and the effects on average recruiting and discharge rates of implementing the legacy policy would likely be best done through a pilot study. Reimbursement rates could then also be adjusted by Medicare to make the policy revenue neutral.

Although the legacy policy eliminates the first-order nonstationary recruiting/discharge behavior (as captured by the fluid model), there may still be a second-order nonstationary behavior due to the uncertainty in patient arrivals and LOS realizations. More specifically, the hospice manager may still have an incentive to recruit cancer patients toward the end of the Medicare year if she is cap constrained as a result of inherent uncertainty in the environment. To test the magnitude of such effects, we again used simulation with a heuristic, analogous to that in §4.3, designed to maximize expected revenue minus costs under the proposed new accounting system (see Online Appendix C for details).

Figure 5 shows the various possibilities for recruiting type 1 patients under the legacy and original heuristics. Note that the recruiting levels under the legacy heuristic are remarkably stationary until the last month or so of the year, whereas the recruiting levels under the original heuristic vary widely throughout the year. We found similar results for the recruitment of type 2 patients and discharges of type 1 patients, although discharges of type 2 patients were only slightly more stationary under the legacy policy, possibly because its decreasing failure rate distribution is being approximated by its mean in our heuristic. This can be seen in Table 2, which shows the time-average squared coefficient of variation of the recruiting and discharge rates across the simulated sample paths. It appears likely that the counteraction to stochastic variability at the end of year is less concerning to Medicare than a calculated nonstationarity resulting from the current policy's incentives.

Although more complicated policy structures could be designed, there is benefit derived from the relative simplicity of the existing and legacy policies from

Figure 5 **Comparison of Stationarity of the Type 1 Patient Recruiting Rates Between the Heuristic Legacy Policy and the Heuristic Fluid Policy (from §4.3) Under Stochasticity**

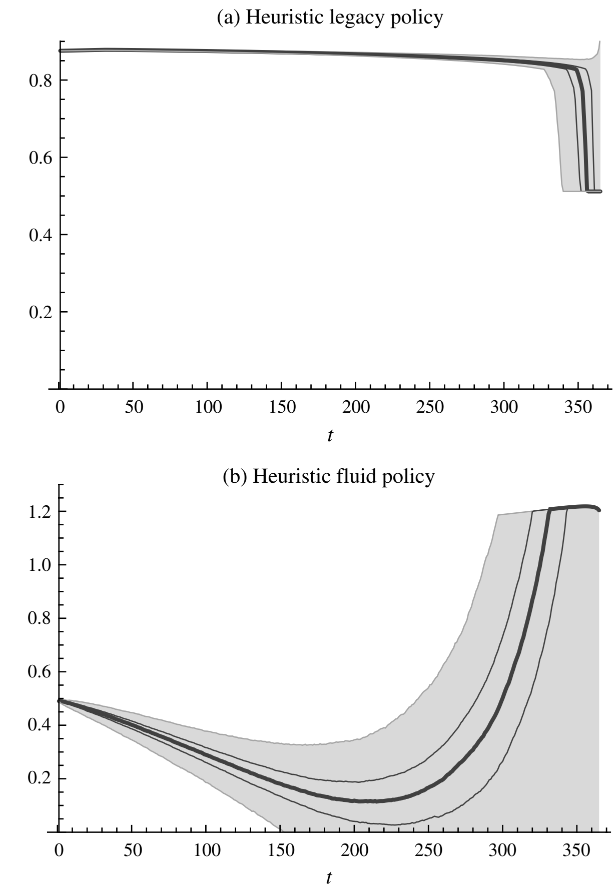

Notes. The bold center line is the median recruiting rate for type 1 patients. The neighboring lines below and above the median line are the 25th and 75th percentile lines; the outside lines are the 1st and 99th percentile lines.

an implementation perspective. Even for these simple policies, an opportunity has grown for Regional Home Health and Hospice Intermediary firms to act as Medicare billing agencies on behalf of the hospices. If the policies became more complicated, then the likelihood is that such intermediaries will consume even more channel profits. A benefit of the legacy policy is that it strongly resembles the existing policy and is unlikely to involve much more difficulty in implementation. The difference is that the hospice will need to segregate any revenues received on behalf of those beneficiaries living into the following year from those newly admitted in that following year.

Table 2 **Average Squared Coefficient of Variation of Recruiting and Discharge Rates Across the Simulated Sample Paths** ($\psi = 0.15$, $R(0) = 55{,}000{,}000$)

	Recruit		Discharge	
	Type 1	Type 2	Type 1	Type 2
Base case	0.766	0.597	1.949	0.242
Legacy	0.005	0.070	0.059	0.154

6. Concluding Remarks and Discussion

We have examined Medicare's hospice reimbursement policy from both profitability and patient recruitment/live-discharge perspectives. Our primary goal has been to inform policy makers and stir debate. From the profitability perspective, we discerned the aspects of the policy or market conditions that lead to potential losses. Specifically, we find that if a provider has a lack of scale and/or an imbalance in the mix of patients, they run a risk of not receiving sufficient revenues to reach profitability. This could be because the patient census lived too long and the cap limited the revenues gained from Medicare while the provider continued to incur the costs of caring for the patients, or because the patient census did not live long enough for the provider to gain sufficient revenues to cover the fixed overhead of operation. We suggest that the government encourage the merging of *appropriate* providers or, at a minimum, remove regulatory hurdles that deter such mergers as a potential remedy.

Another aspect of the current Medicare hospice reimbursement policy we investigated was the manager's optimal recruitment and live-discharge policies. Using a dynamic fluid model, we demonstrate that the manager has an incentive to recruit patients with different diseases at rates that differ according to disease type and that change during the Medicare year. For example, the manager might seek to recruit type 1 (short LOS) patients at a rate that dominates the recruitment rate for type 2 (long LOS) patients toward the end of the year. The basic reason for this is that type 1 patients can increase the cap by the same amount as type 2 patients but are expected to live for shorter durations, which is important when patients living into the next Medicare year are taken into consideration. We also show that the manager has differing incentives for discharging living patients whose conditions have stabilized throughout the year. These are clearly unintended and disturbing consequences of the current policy, and there is strong anecdotal evidence to suggest that such behavior occurs in practice (e.g., Jenkins et al. 2011).

Medicare (and Medicaid) funds a variety of programs aside from hospices. A natural instinct of the policy maker when constructing a set of rules that distributes taxpayer money is to protect the reserves from people trying to take advantage of the system. This is the motivation for the cap. Other public programs have similar caps to limit the government's exposure. The specifics of each program may differ somewhat, but we believe that our lessons on the unintended consequences of such caps are likely to be transferable and that our models can act as prototypes for anyone wishing to model the specifics of those programs.

We design and analyze an alternative policy, called the *legacy policy*, which allows the hospice to continue receiving revenues for these remnant patients, providing any positive remnant cap exists at the end of the year, until the last patient expires or the remnant cap is exhausted, whichever occurs first. It is important that the remnant cap and the remnant patients are tracked separately by the new cap launched at the start of the new year and any newly admitted patients. We show that this alternative policy restores stationarity to the manager's problem (at least in the deterministic model studied), which is compatible with an objective of equal access to hospice care. An attractive attribute of the alternative policy is that it closely resembles the existing policy so that its implementation is not expected to be disruptive to Medicare or hospice providers.

Even with the legacy policy in place, however, the optimal recruitment and discharge rates for different disease types will differ. There are a number of possible remedies for this. Medicare could reimburse at different rates for different diseases, they could adjust the cap increment for different disease types, or they could move to a fixed-plus-variable reimbursement for differing patients. All of these suggestions raise significant new issues, such as the classification of patients with comorbidities into a single class, the incentives for patient churn inherent in a fixed payment system, and the difficulty of calibrating payment rates to actual costs. Our focus has been on highlighting and alleviating the calendar-based recruitment incentives inherent in the current policy. We leave a broader policy study of all incentives under the current scheme as the subject of future research.

To summarize, under any government policy there will always be unintended consequences. This paper sheds light on some of those under the Medicare hospice reimbursement program. In particular, we studied both the efficacy of the program with respect to hospice profitability and the hospice providers' incentives for patient recruitment and live discharges under the program. The primary remedies we suggest—namely, the merging of appropriate providers and the new legacy policy—seek to mitigate the consequences of an undesirable patient mix at a hospice. With respect to implementing the legacy policy, Congress may wish to run a pilot program with the new policy or task the Center for Medicare and Medicaid Studies to use historical data to estimate its financial impact.

Acknowledgments

The authors thank department editor Christian Terwiesch, the associate editor, and two anonymous referees, whose efforts have considerably improved the paper. The authors also thank Nicholas Christakis (Harvard University) and Jack Iwashyna (University of Michigan) for providing patient longevity data, Robert Town (University of Pennsylvania) for valuable comments, and Laurie Meneades for the data transfer. Financial support from the University of Chicago Booth School of Business is gratefully acknowledged.

References

Brandeau ML, Sainfort F, Pierskalla WP, eds. (2004) *Operations Research and Health Care: A Handbook of Methods and Applications* (Kluwer Academic, Norwell, MA).

Campbell DE, Lynn J, Louis TA, Shugarman LR (2004) Medicare program expenditures associated with hospice use. *Ann. Internal Medicine* 140(4):269–277.

Carlson MDA, Herrin J, Du Q, Epstein AJ, Cherlin E, Morrison RS, Bradley EH (2009) Hospice characteristics and the disenrollment of patients with cancer. *Health Services Res.* 44(6):2004–2021.

Centers for Medicare and Medicaid Services (2009) Hospice data, 1998–2008. Report, Department of Health and Human Services, Washington, DC.

Centers for Medicare and Medicaid Services (2010) Processing hospice claims. *Medicare Claims Processing Manual*, Chap. 11 (Revision 1955) (Department of Health and Human Services, Washington, DC).

Christakis NA, Escarce JJ (1996) Survival of Medicare patients after enrollment in hospice programs. *New England J. Medicine* 335(3):172–178.

Eppen GD (1979) Effects of centralization on expected costs in a multi-location newsboy problem. *Management Sci.* 25(5):498–501.

Evans RV (1967) Inventory control of a multiproduct system with a limited production resource. *Naval Res. Logist. Quart.* 14(2):173–184.

Fraser I (1985) Medicare reimbursement for hospice care: Ethical and policy implications of cost-containment strategies. *J. Health Politics, Policy Law* 10(3):565–578.

Huskamp HA, Buntin MB, Wang V, Newhouse JP (2001) Providing care at the end of life: Do Medicare rules impede good care? *Health Affairs* 20(3):204–211.

Jenkins TM, Chapman KL, Ritchie CS, Arnett DK, McGwin G, Cofield SS, Maetz HM (2011) Barriers to hospice care in Alabama: Provider-based perceptions. *Amer. J. Hospice Palliative Medicine* 28(3):153–160.

Killaly BL, Mukamel DB (2010) The cost structure of hospice providers. Working paper, University of California, Irvine, Irvine.

Kutner JS, Meyer SA, Beaty BL, Kassner CT, Nowels DE, Beehler C (2004) Outcomes and characteristics of patients discharged alive from hospice. *J. Amer. Geriatrics Soc.* 52(8):1337–1342.

MedPAC (Medicare Payment Advisory Commission) (2006) Testimony: Report to the Congress: Medicare Payment Policy, U.S. House of Representatives, Committee on Ways and Means, Subcommittee on Health, Washington, DC.

MedPAC (Medicare Payment Advisory Commission) (2010) Testimony: Report to the Congress: Medicare Payment Policy, U.S. House of Representatives, Committee on Ways and Means, Subcommittee on Health, Washington, DC.

Pyenson B, Connor S, Fitch K, Kinzbrunner B (2004) Medicare cost in matched hospice and non-hospice cohorts. *J. Pain Symptom Management* 28(3):200–210.

Rockafellar RT (1970) Conjugate convex functions in optimal control and the calculus of variations. *J. Math. Anal. Appl.* 32(1):174–222.

Sack K (2007) In hospice care, longer lives mean money lost. *New York Times* (November 27) http://www.nytimes.com/2007/11/27/us/27hospice.html.

Taylor DH Jr, Ostermann J, Van Houtven CH, Tulsky JA, Steinhauser K (2007) What length of hospice use maximizes reduction in medical expenditures near death in the U.S. Medicare program? *Soc. Sci. Medicine* 65(7):1466–1478.

Taylor DH Jr, Steinhauser K, Tulsky JA, Rattliff J, Van Houtven CH (2008) Characterising hospice discharge patterns in a nationally representative sample of the elderly, 1993–2000. *Amer. J. Hospice Palliative Medicine* 25(1):9–15.

U.S. Government Accountability Office (2004) Medicare hospice care: Modifications to payment methodology may be warranted. Report to the Committee on Finance, U.S. Senate, and the Committee on Ways and Means, House of Representatives, Washington, DC.

Wright AA, Katz IT (2007) Letting go of the rope—Aggressive treatment, hospice care, and open access. *New England J. Medicine* 357(4):324–327.

Zipkin PH (2000) *Foundations of Inventory Management* (McGraw-Hill, New York).