

Does hospice reduce costs at the end of life? Estimates from a dynamic model of hospice choice

Toby C. Campbell
University of Wisconsin
tcc@medicine.wisc.edu

Andrea P. Chung
The Lewin Group
park.andrea@gmail.com

Dennis McWeeny
University of Wisconsin
dmcweeny@wisc.edu

Alan T. Sorensen*
University of Wisconsin & NBER
sorensen@ssc.wisc.edu

October 2018

PRELIMINARY DRAFT: PLEASE DO NOT CIRCULATE OR CITE

Abstract

We use Medicare claims data to measure how hospice enrollment affects cancer patients' costs and mortality. We directly model the endogeneity of the hospice enrollment decision, framing patients' choices of whether and when to enroll as a single-agent dynamic discrete choice problem with an unobserved state variable (health status). Preliminary estimates of the model indicate that the option to enroll in hospice leads to substantial cost savings, and almost all of these savings result from lower daily costs rather than accelerated mortality.

*This project was supported by the Health Innovation Program, the UW School of Medicine and Public Health from the Wisconsin Partnership Program, and the Community-Academic Partnerships core of the University of Wisconsin Institute for Clinical and Translational Research through the National Center for Advancing Translational Sciences, grant UL1TR000427. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH.

1 Introduction

End-of-life care figures prominently in policy debates about how to control health care costs. Medicare spends over \$200 billion each year—more than a quarter of its total expenditures—on care for patients in their last year of life, and annual per capita spending on inpatient hospital services is roughly seven times higher for patients who died than those who did not. To some extent these cost differences are unavoidable, since dying patients are naturally sicker and more expensive to treat. But to the extent that the cost differences reflect overuse of resource-intensive treatments at the end of life, limiting such treatments could significantly reduce overall costs without meaningfully reducing patients’ quality of life—especially given that some studies have suggested that higher costs are associated with a worse quality of death (Zhang et al (2009)).

Hospice has played an increasingly important role in end-of-life care, as the share of dying patients who enroll in hospice grew from 25 percent in 2000 to 46 percent in 2015.¹ Hospice provides palliative care, such as pain management and emotional counseling, in lieu of aggressive medical intervention. Since patients who choose hospice are choosing to forego expensive curative-treatments, hospice has the potential to substantially reduce Medicare’s costs, both because hospice patients have lower daily costs (since palliative care is cheaper than curative care) and because they may live fewer days (as a consequence of foregoing life-prolonging treatments).

Our aim in this paper is to use Medicare claims data to directly measure how hospice enrollment affects cancer patients’ costs and mortality. That is, if a patient chooses hospice, how much does this reduce the expected costs of that patient’s end-of-life care? And how much of the cost reduction is due to accelerated mortality? These effects are difficult to measure empirically, due to the endogeneity of the hospice enrollment decision. Patients who choose hospice tend to be sicker than non-hospice patients in ways the econometrician cannot observe, so simple comparisons of hospice vs. non-hospice patients’ costs and mortality would obviously be misleading. Moreover, hospice enrollment is not only an endogenous decision, it is a dynamic one: one cannot simply compare the outcomes for patients who did vs. did not choose hospice, since this would ignore the arguably more important decision of *when* to enter hospice.

To address these challenges, we develop and estimate a model that frames patients’ hospice enrollment decisions as a single-agent dynamic discrete choice problem with an unobserved state variable. In the model, a patient who is diagnosed with metastatic cancer begins with some initial unobserved health state $\theta_0 > 0$, and this health state then declines stochastically over time until it crosses zero, at which point the patient dies. The hospice enrollment decision is driven by differences in flow utilities, which depend on the patient’s health state. If a patient enrolls in hospice, she earns a constant flow utility until she dies. While not enrolled in hospice, her flow utility is an increasing function of her health state, and this flow utility goes to zero as her health state approaches zero

¹See Han et al (2006) and NHPCO (2016).

(i.e., as she approaches death). Thus, the patient enrolls in hospice once her health state falls below some optimal threshold. This threshold is partly a function of a patient’s idiosyncratic preferences for hospice, but we also assume it can be shifted by instruments that are exogenous to the patient’s health state. We use as our primary instrument the overall hospice utilization rate in the patient’s hospital service area in the year prior to the patient’s diagnosis. Thus, our structural model makes explicit the endogeneity of the hospice enrollment choice, and also exploits plausibly exogenous variation in the data to identify the model’s key parameters.

Our empirical analysis focuses on a sample of over 30,000 patients who were diagnosed with metastatic cancer between 2005 and 2011. Among the patients who died before the end of our sample period, 61% enrolled in hospice before dying. Unconditional average costs are lower for hospice patients, even though their average survival times are longer than for non-hospice patients. We show that average daily costs increase significantly in the days just prior to death, but this increase is not as sharp among hospice patients. We also show that patients in regions with high hospice utilization rates tend to enroll in hospice slightly sooner and also die slightly sooner.

Preliminary estimates of the model indicate that the option to enroll in hospice leads to substantial cost savings, and almost all of these savings result from lower daily costs rather than accelerated mortality. Our model does imply that forgoing curative care stochastically shortens one’s life, but the effect is small, accounting for less than 5 percent of the overall cost savings from hospice.

2 A brief background on hospice care

Hospice takes a palliative approach to end-of-life care, aiming to allow patients to die “pain-free and with dignity.” Medicare began coverage of hospice services in 1983, and since then there has been an increasing acceptance of hospice as a viable mode of care for the terminally ill. In the U.S. in 2015 there were 4,199 hospice providers that served 1.38 million patients. Nearly half of Medicare decedents were enrolled in hospice at the time of death.²

Patients usually enroll in hospices through referrals from physicians and long-term care facilities. In order to be eligible for Medicare’s hospice benefit, the patient’s physician and the hospice medical director must certify that the patient has a terminal disease with a life expectancy of 6 months or less. Hospice services are typically provided in the home. A family member usually serves as the primary caregiver, while hospice staff members—including physicians, nurses, home health aides, social workers, clergy and other volunteers—make regular visits to the patient and provide other support services for the family. Typical hospice services include pain management, emotional/psychosocial counseling, drugs and medical equipment, and bereavement care for family and friends.

²See National Hospice and Palliative Care Organization: Facts and Figures, 2016.

Medicare reimbursement for hospice services is on a per diem basis, with four tiers of reimbursement rates corresponding to different intensities of care. The payment rates were set based on data from a Medicare demonstration project done in the early 1980s, and are adjusted annually to reflect inflation. The vast majority of patient days are categorized as Routine Home Care, for which the reimbursement rate is roughly \$150 per day. During brief periods of crisis (e.g., to manage acute medical symptoms), the reimbursement level can be escalated to Continuous Home Care, which is reimbursed at roughly \$900 per day. When patients need pain control or other symptom management that cannot be managed at the patient’s home, they can be hospitalized on a short-term basis in a hospice inpatient facility, or a hospital with which the hospice has contracted, and the hospice is reimbursed for General Inpatient Care at roughly \$700 per day. Finally, when inpatient care is needed on a short-term basis in order to relieve the beneficiary’s primary caregiver, the hospice can be reimbursed for Inpatient Respite Care at roughly \$160 per day. In addition to the annual inflation updates, these reimbursement rates are also adjusted locally using area wage indexes to reflect differential labor costs in different markets. The per-diem rates include payment for drugs used for pain control and symptom management; drugs prescribed for conditions unrelated to the patient’s terminal illness must be reimbursed separately.

3 Previous studies of hospice costs

Given the growing importance of hospice and the strong a priori arguments for potential cost savings, it is not surprising that several previous studies have attempted to measure cost differences between hospice and non-hospice patients. These studies are best characterized by how they address two fundamental challenges: endogeneity of the hospice choice, and heterogeneity in the timing of hospice enrollment. One thread in the literature addresses the endogeneity problem with propensity score methods. For example, [Taylor et al \(2007\)](#) study a random sample of Medicare decedents who died between 1993 and 2003, and use propensity score matching to match hospice patients to a control group of patients who did not enter hospice but had similar characteristics. Comparing the costs of hospice patients to the costs of the matched control patients, they find that hospice patients’ costs were on average \$2,309 lower in the period after hospice initiation. However, total costs in the year of death were not statistically different for hospice patients when compared to non-hospice patients (\$32,727 vs. \$33,837) because the average costs for hospice patients *prior to hospice enrollment* were higher for the hospice group. [Kelley et al \(2013\)](#) take a similar approach, but improve on the propensity score matching by augmenting Medicare claims with matched data from the Health and Retirement Study, which includes information about patients’ functional status and social characteristics. They find cost reductions of roughly the same magnitude as the Taylor et al (2007) study, and also document reductions in hospital days and ICU days for patients who choose hospice.

A drawback of these studies is that they employ what Kelley *et al* (2013) call a “mortality follow-

back” design. To compare the costs incurred by a hospice patient vs. non-hospice patient, they consider costs in the last N weeks of life for the two patients, counting back from the date of death. While useful, such comparisons do not capture any differences in the length of life for the two patients—and such differences may be important, since hospice patients do not typically receive curative treatments. An alternative approach that does allow for these differences is what might be called a “follow-forward” design. For example, [Pyenson et al \(2004\)](#) and [Connor et al \(2007\)](#) compare costs and times until death for terminally ill hospice and non-hospice Medicare patients within narrowly defined trigger diagnoses—i.e., diagnoses identified by a group of physicians as indicators of hospice eligibility. The researchers calculated all costs between time of diagnosis and time of death for each patient with one of these trigger diagnoses, and then compared the costs for patients who chose hospice to the costs of those who did not. [Pyenson et al \(2004\)](#) found lower mean costs for hospice patients overall, but for most diagnoses the differences were not statistically significant. Both the [Pyenson et al \(2004\)](#) and [Connor et al \(2007\)](#) studies report that patients who chose hospice actually lived longer on average than patients who did not; they suggest this may simply reflect that patients who live longer have more time to choose hospice. While these studies have the advantage of being able to analyze mortality differences for hospice vs. non-hospice patients, the main drawback is that they ignore any patient heterogeneity within diagnosis categories—that is, they do not attempt to control for the possibility that among patients with the same trigger diagnosis, the sickest of these patients may be the ones who choose hospice. To some extent, differences in patients’ health status can be controlled for with observable covariates,³ but unobservable patient heterogeneity is also undoubtedly important. Additionally, while the setup of these “follow-forward” studies is conceptually appealing, it turns out that patients who choose hospice often wait several weeks or months before doing so. This means that simply comparing hospice patients to non-hospice patients ignores what appears to be an important decision margin: *when* to enroll in hospice.

Our study aims to incorporate the desirable aspects of both the aforementioned types of studies. We analyze our sample of patients prospectively, using a “follow-forward” rather than “follow-back” design, so we can measure any effects of hospice on mortality. Instead of simply comparing hospice patients to non-hospice patients, we develop a dynamic discrete choice model of patient’s decisions about when to enter hospice. To address endogeneity, our model explicitly accounts for unobserved health status, and in estimating the model we exploit instrumental variables that generate plausibly exogenous variation in patients’ propensity to choose hospice. In this respect, our paper is similar to others that have modeled dynamic decision-making in the presence of an unobserved (to the researcher) state variable. For example, [Fang and Kung \(2012\)](#) use a dynamic discrete choice model with serially correlated unobserved state variables to empirically analyze individuals’ decisions to let their life insurance policies lapse. [Chan and Hamilton \(2006\)](#) develop a dynamic model of patients’ decisions to drop out of a clinical trial for HIV drugs; in their model,

³See, for example, [Kelley et al \(2011\)](#), which shows that observable covariates meaningfully improve predictions of Medicare expenditures in the last six months of life.

the patient-specific side effects of a drug are the key unobservable factor.

4 Data

Our data consist of Medicare claims for a sample of 30,102 patients who were diagnosed with metastatic cancer between 2005 and 2011. The patients were identified and selected from a 5% random sample of all Medicare patients by finding patients with either (a) a “bad” cancer that had metastacized; or (b) any cancer that had metastacized to the brain or liver.⁴ Prognosis for such patients is generally grim. Measuring from the date of the first appearance of a secondary neoplasm in the claims data, 57 percent of the patients in our sample died within six months, and 69 percent died within a year. Table 1 lists the different cancer types and the corresponding mortality rates.

Table 1: Mortality rates by cancer type

| Primary cancer type | # patients | Percent dying within: | | |
|---------------------|------------|-----------------------|----------|-----------|
| | | 3 months | 6 months | 12 months |
| Esophagus | 629 | 0.41 | 0.55 | 0.66 |
| Stomach | 1,099 | 0.36 | 0.49 | 0.63 |
| Liver | 2,493 | 0.48 | 0.60 | 0.73 |
| Gallbladder | 554 | 0.46 | 0.60 | 0.75 |
| Pancreas | 2,102 | 0.53 | 0.65 | 0.77 |
| Lung | 11,783 | 0.44 | 0.56 | 0.69 |
| Pleura | 804 | 0.38 | 0.52 | 0.67 |
| Brain | 1,405 | 0.34 | 0.51 | 0.65 |
| Other | 9,233 | 0.47 | 0.58 | 0.68 |
| Total | 30,102 | 0.45 | 0.57 | 0.69 |

Survival times are calculated relative to the date of the first appearance of metastasis, as indicated by a claim with an ICD-9 code for a secondary neoplasm. To avoid censoring, patients whose initial diagnosis occurred within 12 months of the end of our data sample were excluded from the calculations in this table. Patients with cancer type “Other” are patients with primary cancers in other sites where the metastases were to the brain or liver.

The claims data cover the period 2004-2011 and include Medicare parts A (hospital, skilled nursing facility, and hospice claims), B (physician services, lab tests, and imaging claims), and D (prescription drug claims).⁵ For each claim we know the date, claim type, HCPCS or CPT code describing the procedure or service, and cost (amount paid by Medicare). We determine the date of hospice enrollment as the first date on which a hospice claim is observed.

⁴The specific criteria for inclusion were either (a) having a claim with an ICD-9 code of 150 (Malignant Neoplasm of Esophagus), 151 (Stomach), 155 (Liver), 156 (Gallbladder), 157 (Pancreas), 162 (Lung), 162 (Pleura), or 191 (Brain), and also a claim with an ICD-9 code of 196, 197, or 198 (Secondary malignant neoplasms); or (b) having a claim with an ICD-9 code for any cancer (140-209), and also a claim with an ICD-9 code of 198.3 (Secondary malignant neoplasm of brain) or 197.7 (Secondary malignant neoplasm of liver).

⁵Since the Medicare prescription drug program began in 2006, Part D claims are only available from 2006 on.

Table 2 reports summary statistics describing the sample patients’ characteristics. Since these are Medicare-eligible patients, the average age is relatively high. The HRR hospice utilization rate is the average hospice utilization rate in the patient’s home hospital referral region, calculated in the year prior to the patient’s initial metastatic diagnosis. The HCC score is a CMS risk score based on the patient’s previous claims, computed at the time of the initial metastatic diagnosis; higher scores indicate sicker patients.⁶ In our sample, patients with HCC scores below the mean (2.07) had a 6-month mortality rate of 55 percent, while the rate for patients with HCC scores above the mean was 61 percent. Of course, all of the patients in our data were quite ill, so the costs of treating them were high: from the date of diagnosis of metastatic cancer to the date of death, the average total cost per patient was \$58,378.

Table 2: Patient Characteristics

| | Mean | Std. Dev. | Percentiles | | |
|----------------------|--------|-----------|-------------|--------|---------|
| | | | 0.05 | 0.50 | 0.95 |
| Age | 77.69 | 7.09 | 67.28 | 77.20 | 90.02 |
| Female | 0.54 | 0.50 | 0.00 | 1.00 | 1.00 |
| HCC score | 2.07 | 1.40 | 0.44 | 1.78 | 4.79 |
| HRR hospice util. | 0.45 | 0.21 | 0.11 | 0.44 | 0.86 |
| Total inpatient days | 16.26 | 19.36 | 0.00 | 11.00 | 51.00 |
| Total hospice days | 116.82 | 497.65 | 0 | 0 | 564 |
| Total costs | 58,378 | 54,740 | 8,078 | 42,528 | 163,473 |
| Average daily costs | 1,072 | 3,297 | 37 | 362 | 3,708 |

Age and HCC score are recorded on the date of the first diagnosis of metastatic cancer. HRR hospice utilization is the average hospice utilization rate (hospice deaths divided by total deaths) in the patient’s hospital referral region. Total inpatient days, hospice days, and costs are calculated from the date of metastatic cancer diagnosis to the date of death. Censored patients—those who were still alive at the end of our sample period—were excluded from the calculations for inpatient days, hospice days, and costs.

4.1 Patterns in Hospice Utilization, Costs, and Mortality

Among the 22,474 patients in the sample who died before the end of the sample period, 60.7 percent enrolled in hospice before dying. Table 3 shows a simple comparison of unconditional mean costs and survival times for patients who chose hospice vs. those who did not. Total costs—i.e., total charges to Medicare from the time metastasis was diagnosed until death or censoring at the end of the sample period—were roughly \$2,400 lower for hospice patients, and inpatient costs were roughly \$4,900 lower.

Costs were lower for hospice patients even though on average they lived longer than non-hospice patients. As noted by Pyenson et al (2004), the longer survival times for hospice patients may

⁶See Pope et al (2000) for a detailed explanation of the CMS Hierarchical Condition Categories risk scoring protocol.

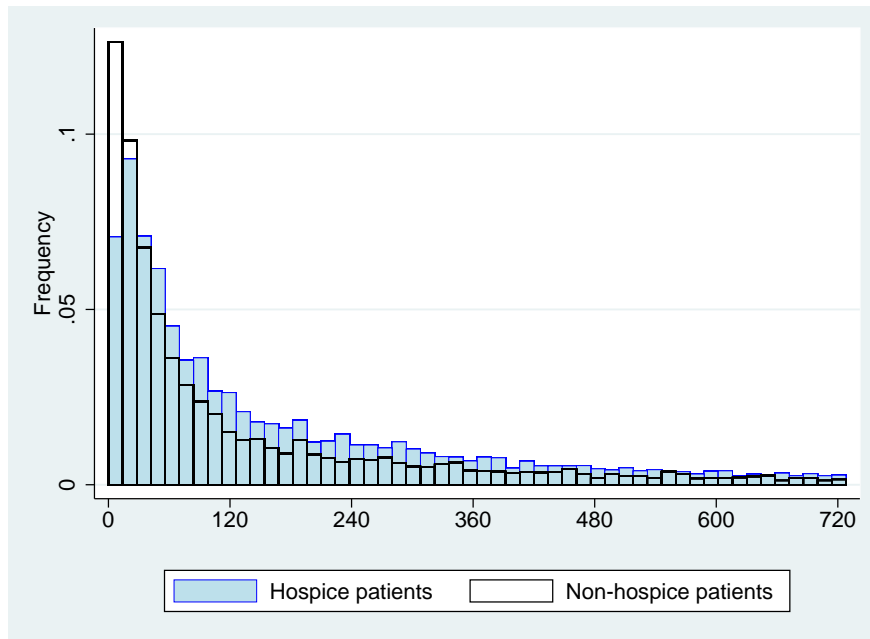
Table 3: Differences in costs and survival times

| | Patients who enrolled in hospice | Patients who never enrolled in hospice |
|----------------------|-------------------------------------|---|
| Total costs | 46,292 (53,801) | 48,738 (58,163) |
| Inpatient costs | 15,012 (20,515) | 19,891 (26,506) |
| Survival time (days) | 257 (355) | 193 (317) |

Cells report means and standard deviations (in parentheses). Patients whose initial diagnosis was in 2010 or 2011 are excluded from these calculations, in order to minimize the potential impact of censoring. Total costs are the sum of all claims from diagnosis to death; inpatient costs are the sum of all inpatient costs.

reflect the fact that patients who live longer have more of an opportunity to choose hospice. For example, patients who die shortly after the diagnosis of metastasis may not have had time to make the decision and transition to hospice. However, while the full distribution of survival times (shown in Figure 1) does indicate that non-hospice patients were more likely to die within a few days of diagnosis, the shapes of the distributions are otherwise similar, and there is a thick tail of long survival times for hospice patients.

Figure 1: Distributions of survival times



Simple comparisons between patients who did vs. did not enroll in hospice ignore a dimension of the data that appears to be important: the timing of the hospice enrollment decision. Conditional on enrolling in hospice, there is wide variation in how long it takes patients to do so. Approximately 25% of patients who choose hospice enroll within two weeks after the diagnosis of metastasis, but another 25% wait six months or longer before enrolling. Presumably much of this variation reflects heterogeneity in the severity of illness, though we expect some of it reflects heterogeneity in patients’ preferences as well. There is also substantial variation across patients in the time between hospice enrollment and death: roughly half of hospice patients die within two weeks after enrollment, but over 10% live for three months or longer after enrolling.

The model we develop below is focused on three time intervals: the time between diagnosis and hospice enrollment, the time between hospice enrollment and death, and the overall time between diagnosis and death (which, for hospice patients, is simply the sum of the first two intervals). Table 4 shows the average lengths of these intervals for patients who chose hospice, in regions with high vs. low hospice utilization. The numbers suggest that hospice enrollment decisions are influenced by location. On average, patients in high-utilization regions enroll in hospice roughly 4 days sooner than patients in low-utilization regions. If enrolling in hospice—and therefore foregoing curative treatments—had no impact on mortality, then the accelerated hospice enrollments of patients in high-utilization regions would not lead to accelerated mortality; but the average overall survival times are 3 days shorter for patients in high-utilization regions. We return to this point when discussion identification of our model in Section 5.1.

Table 4: Hospice enrollment and mortality in high- vs. low-utilization regions

| | Days between: | | |
|--------------------------|----------------------------------|------------------------------|---------------------|
| | Diagnosis and hospice enrollment | Hospice enrollment and death | Diagnosis and death |
| Low-utilization regions | 126.6 (158.0) | 36.2 (61.2) | 162.8 (169.9) |
| High-utilization regions | 122.3 (158.8) | 37.6 (67.4) | 159.8 (171.8) |
| Overall | 124.3 (158.4) | 36.9 (64.6) | 161.2 (170.9) |

Cells report means and standard deviations (in parentheses). The sample used for these calculations included only patients who chose hospice, and whose initial diagnosis came before 2010. Patients’ home regions (hospital service areas) are categorized as high-utilization regions if the overall hospice utilization rate was higher than the median (across patients) in the relevant year.

Table 5 describes the distributions of daily costs for hospice days and non-hospice days. A very

small fraction of observations report negative costs or extremely high costs, presumably because a provider was making a correction to previous submitted claims, or submitting a single claim for services that spanned several days. To mitigate the influence of these anomalies, we trimmed the sample by excluding observations with reported daily costs below the 0.25th or above the 99.75th percentiles.


Table 5: Daily costs

| | Non-hospice days | Hospice days |
|------------------------|------------------|--------------|
| Number of patient-days | 8,792,869 | 737,390 |
| Average costs | 130.04 | 176.12 |
| Standard deviation | 442.13 | 258.87 |
| Percentiles: | | |
| 0.05 | 0.00 | 0.00 |
| 0.25 | 0.00 | 108.23 |
| 0.50 | 0.00 | 135.36 |
| 0.75 | 30.15 | 157.60 |
| 0.95 | 830.53 | 563.27 |

Days with reported costs below the 0.25th or above the 99.75th percentile were excluded from these calculations.

Figure 2 illustrates how daily costs change at the end of life. The figure shows smoothed plots of average daily costs, estimated using local polynomial regression, as a function of time to death. At 120 days prior to death, average daily costs are similar for hospice patients and non-hospice patients. Unsurprisingly, costs increase as the patient approaches death, and the increase is much more pronounced for non-hospice patients. The figure thus illustrates what we might expect to be the principal effect of hospice: eliminating expensive medical interventions at the very end of life.

Figure 2: Costs at the end of life



5 Model

Our overall objective is to build a model that explains patients’ decisions about hospice, as well as the implications of those decisions for costs. We begin by outlining our model of the hospice decision. Since our model of costs is straightforward—we simply parameterize distributions of daily costs for hospice and non-hospice patients—we postpone its description to section 5.1, where we discuss the details of our estimation procedure.

We model the patient’s hospice enrollment decision as a dynamic choice problem in continuous time. The patient chooses whether to enroll in hospice at any time t , indicated by $h_t \in \{0, 1\}$. This choice depends on the patient’s health state, θ_t , which is fully known to the patient but unobserved by the econometrician. This health state starts at some initial value $\theta_0 > 0$ (at the time of the first diagnosis of metastasis), and then declines stochastically until it falls below zero, at which point the patient dies. Formally, we model θ_t as a Brownian motion with drift:

$$d\theta_t = \mu dt + \sigma dW_t , \tag{1}$$

where dW_t is a standard Wiener process and the drift is

$$\mu = \begin{cases} -1 & \text{if } h_t = 1 \\ -(1 + \kappa) & \text{if } h_t = 0 \end{cases} \quad (2)$$

The parameter κ captures any effects of hospice enrollment on mortality. If $\kappa < 0$, then health declines less rapidly for non-hospice patients—e.g., because they are receiving life-prolonging treatments—so choosing hospice means accepting a shorter expected survival time.

The initial value θ_0 , which represents the patient’s health state when he is first diagnosed with metastatic cancer, is assumed to be a function of patient characteristics:

$$\theta_0 = \bar{\theta} + x' \alpha \quad (3)$$

where x is a vector of characteristics such as the patient’s sex, age at diagnosis, and HCC score at diagnosis.

Patients trade off the costs of hospice (accelerated mortality) against the benefits of higher flow utility. When health deteriorates sufficiently—i.e., for low enough values of θ_t —some patients will prefer to enroll in hospice even if it stochastically shortens their lives. This is captured in our model by differences in flow utility, which is given by

$$u(h_t; \theta_t) = \begin{cases} \theta_t & \text{if } h_t = 0 \text{ and } \theta_t > 0 \\ \Delta + z' \gamma + \epsilon & \text{if } h_t = 1 \text{ and } \theta_t > 0 \\ 0 & \text{if } \theta_t \leq 0 \end{cases} \quad (4)$$

For non-hospice patients, flow utility is simply the health state θ_t , which captures the notion that better health results in a higher quality of life, and also builds into the model the idea that the healthier the patient, the less likely she will be to enroll in hospice.

Flow utility while on hospice depends on three factors. The parameter Δ represents the average utility, common across patients, that results from being on hospice—for example due to palliative treatments or the value of being at home vs. in the hospital. The instruments z are observable characteristics of the patient or the patient’s market that shift the perceived costs or benefits of hospice; in our application, the main instrument is the previous year’s hospice utilization rate in the patient’s HSA. The ϵ term is an unobservable taste shock, assumed to be *iid* across patients, intended to capture heterogeneity in preferences for hospice. We assume that ϵ is a draw from a normal distribution with mean zero and standard deviation σ_ϵ , and that ϵ ’s are *iid* across patients.

We assume that, from the patient’s perspective, enrolling in hospice is an irreversible decision. The value function can then be expressed as

$$V(\theta; z, \epsilon) = \max \left\{ \int_0^\infty \left[\int_0^t e^{-\rho s} h(z, \epsilon) ds \right] f(t|\theta) dt, \theta + \frac{1}{1 + \rho dt} \mathbb{E}[V(\theta + d\theta; z, \epsilon)|\theta] \right\} \quad (5)$$

where $f(t|\theta)$ is the density of time to death conditional on the current health state θ . The solution to the patient's decision problem is a threshold rule: given the patient's z and ϵ , there is some threshold $\bar{\theta}(z, \epsilon)$ such that when θ_t reaches this threshold, the patient will enroll in hospice. For low enough values of ϵ , this threshold may be negative; in that case, the patient will never enroll in hospice. In fact, this is the only way the model can rationalize patients who die without enrolling in hospice, which is a point we will return to below when interpreting our results.

We make one further simplifying assumption, which is that patients do not discount future payoffs: $\rho = 1$. One rationale for this assumption is that since patient's expected survival times are relatively short, time discounting may not be quantitatively important in their decision-making. But the main motivation for this assumption is computational convenience. When $\rho = 1$, we can derive an almost-analytic solution for the threshold θ as a function of z and ϵ : it is the non-trivial solution to the equation⁷

$$\frac{\sigma^2}{2\mu} \left[\exp\left(\frac{2\mu\bar{\theta}}{\sigma^2}\right) - 1 \right] \left[h(z, \epsilon) + \frac{1}{\mu}\bar{\theta} - \frac{\sigma^2}{2\mu^2} \right] - \frac{1}{2\mu}\bar{\theta}^2 - \left[h(z, \epsilon) - \frac{\sigma^2}{2\mu^2} \right] \bar{\theta} = 0 \quad (6)$$

In estimating the model, this allows us to quickly solve the patient's dynamic decision problem given any guess of the parameters.

Another important convenience that results from modeling the health status as a Brownian motion is that first-passage times have known densities. Letting $g(t|\theta_a, \theta_b, \mu)$ denote the density of the time to first crossing of θ_b , starting at θ_a with a drift of μ , then

$$g(t|\theta_a, \theta_b, \mu) = \frac{\theta_a - \theta_b}{\sqrt{2\pi\sigma^2 t^3}} \exp\left[-\frac{(\theta_a - \theta_b - \mu t)^2}{2\sigma^2 t}\right] \quad (7)$$

The corresponding survivor function—the probability that θ_t hasn't yet crossed θ_b at t —is

$$S(t|\theta_a, \theta_b, \mu) = \Phi\left(\frac{\theta_a - \theta_b - \mu t}{\sqrt{\sigma^2 t}}\right) - \exp\left(-\frac{2(\theta_a - \theta_b)\mu}{\sigma^2}\right) \Phi\left(\frac{\theta_a - \theta_b + \mu t}{\sqrt{\sigma^2 t}}\right) \quad (8)$$

where Φ is the CDF of the standard normal distribution.

⁷The derivation is provided in Appendix ??.

5.1 Estimation

Our model has two parts: the model of patients' hospice enrollment decisions, and the model of daily costs conditional on hospice status. For reasons explained below, we estimate these two components in separate steps.

Hospice enrollment

Estimation of the first component is relatively straightforward, since our model of stochastically declining health implies that patients' decisions are threshold rules. The threshold $\bar{\theta}$ determines the timing of hospice enrollment. After a patient enrolls in hospice, the time until death is the time it takes for the health state—which now follows a Brownian motion with drift $\mu = -1$ —to go from $\bar{\theta}$ to 0. Our model thus generates likelihoods for two key dependent variables in the data: T_H , the time from diagnosis to hospice enrollment; and T_D , the time from hospice enrollment to death.

The parameters to be estimated are α , κ , and σ_η , which govern the process of the unobserved health state; and Δ , γ , and σ_ϵ , which determine the utility differential that results from choosing hospice. Let $\Psi \equiv (\alpha, \kappa, \sigma_\eta, \Delta, \gamma, \sigma_\epsilon)'$ be the vector of these parameters. Let D_i equal one if patient i 's death is observed, and let H_i equal one if patient i entered hospice prior to death. In our model, if a patient dies before entering hospice then her threshold $\bar{\theta}$ must have been negative. Under our assumptions outlined above, the likelihood function for patient i can then be written as

$$L_i(\Psi) = \begin{cases} \int_{\{\epsilon: 0 < \bar{\theta}(z_i, \epsilon) < \theta_0(x_i)\}} g(T_{H,i} | \theta_0(x_i), \bar{\theta}(z_i, \epsilon), \mu_0) g(T_{D,i} - T_{H,i} | \bar{\theta}(z_i, \epsilon), 0, \mu_1) \phi(\epsilon/\sigma_\epsilon) d\epsilon & D_i = 1, H_i = 1 \\ \int_{\{\epsilon: \bar{\theta}(z_i, \epsilon) < 0\}} g(T_{D,i} | \theta_0(x_i), 0, \mu_0) \phi(\epsilon/\sigma_\epsilon) d\epsilon & D_i = 1, H_i = 0 \\ \int_{\{\epsilon: 0 < \bar{\theta}(z_i, \epsilon) < \theta_0(x_i)\}} g(T_{H,i} | \theta_0(x_i), \bar{\theta}(z_i, \epsilon), \mu_0) S(T_{D,i} - T_{H,i} | \bar{\theta}(z_i, \epsilon), 0, \mu_1) \phi(\epsilon/\sigma_\epsilon) d\epsilon & D_i = 0, H_i = 1 \\ \int_{\{\epsilon: \bar{\theta}(z_i, \epsilon) < \theta_0(x_i)\}} S(T_{D,i} | \theta_0(x_i), \max\{0, \bar{\theta}(z_i, \epsilon)\}, \mu_0) \phi(\epsilon/\sigma_\epsilon) d\epsilon & D_i = 0, H_i = 0 \end{cases} \quad (9)$$

where $\mu_0 = -(1 + \kappa)$ is the health drift while a patient is not on hospice, $\mu_1 = -1$ is the drift rate while on hospice, and ϕ is the pdf of the standard normal distribution. Note that for censored patients—i.e., patients who are still alive at the end of our sample period— T_D is still defined as the number of days lived beyond the initial diagnosis.

We estimate the parameters Ψ by maximizing the log-likelihood $\sum_i \ln L_i(\Psi)$, evaluating the integrals over ϵ numerically.

Costs

We model the distribution of daily costs (c_t) using the following parametric functions of hospice enrollment (h_t) and health status (θ_t), suppressing i subscripts for convenience:

$$\Pr(c_t|\theta_t, h_t) = \begin{cases} \lambda_h & \text{if } c_t = 0 \text{ and } h_t = 1 \\ (1 - \lambda_h)\phi\left(\frac{\log(c_t) - \beta_h}{\sigma_{c,h}}\right) & \text{if } c_t > 0 \text{ and } h_t = 1 \\ \lambda_{nh} & \text{if } c_t = 0 \text{ and } h_t = 0 \\ (1 - \lambda_{nh})\phi\left(\frac{\log(c_t) - \beta_0 - \beta_1 \log(\theta_t)}{\sigma_{c,nh}}\right) & \text{if } c_t > 0 \text{ and } h_t = 0 \end{cases}. \quad (10)$$

Conditional on θ_t and h_t , costs are independent across days and across patients, and conditional on being nonzero, costs follow a lognormal distribution. This specification captures several important features of our data. First, there are many days during which patients incur zero costs, especially while not enrolled in hospice. We allow costs to be zero with probability λ_h for hospice patients and probability λ_{nh} for nonhospice patients. Second, costs increase significantly in the days just prior to death, especially for nonhospice patients. To fit this feature of the data, we allow log costs to depend on the log of health status for nonhospice patients; as health status approaches zero (i.e., as the patient approaches death), costs approach infinity in our model.⁸

We estimate λ_h and λ_{nh} directly from the data, since they are given by the fraction of days with zero costs for hospice and non-hospice patients, respectively. The estimation of β_h and $\sigma_{c,h}$ is also a straightforward maximum likelihood problem, since it reduces to estimating the mean and variance of a normal distribution. However, estimation of the non-hospice cost function presents several technical challenges. First, the integral over ϵ must be evaluated numerically just as in the first step described above. Second, since θ_t is a serially correlated unobserved state variable, calculating the likelihood involves evaluating a high-dimensional integral over the joint distribution of $\{\theta_t\}_{t=0}^{\min\{T_{H,i}, T_{D,i}\}}$ for each patient. The direct evaluation of this integral is computationally infeasible, so we adopt a simulation approach instead.

Frequency simulation from the Brownian motion with drift is not straightforward because the simulated sequences must be consistent with observed behavior. In particular, for a patient who enters hospice, we know that $\theta_{T_{H,i}} = \bar{\theta}(z_i, \epsilon)$ and that $\theta_t > \bar{\theta}(z_i, \epsilon)$ for every $t < T_{H,i}$. Similarly, for patients who die, we know that $\theta_{T_{D,i}} = 0$ and that $\theta_t > 0$ for every $t < T_{D,i}$. If we simply simulated R sequences of the Brownian motion with drift, it is extremely likely that none of the sequences would satisfy these restrictions. As a result, the likelihood would be undefined. Clearly, we need a better way of drawing sequences of θ_t .

⁸This does not present a problem for estimation; we assume that a patient wakes up with a positive θ_t on the date of death. As a result, $\log(\theta_t)$ is always well-defined in the likelihood function.

A solution to this problem is found in the theory of restricted Brownian meanders. A restricted Brownian meander B_t^{reme} is a stochastic process for which $B_0^{reme} = a$, $B_T^{reme} = b$ for some $T > 0$, and $B_t^{reme} > b$ for every $t < T$. This process can be simulated using the expression

$$B_t^{reme} = b + \sqrt{\left(\frac{a-b}{T}(T-t) + B_t^{br1}\right)^2 + (B_t^{br2})^2 + (B_t^{br3})^2} \quad (11)$$

where B_t^{br1} , B_t^{br2} , and B_t^{br3} , are independent standard Brownian bridges defined by the stochastic differential equation

$$dB_t^{br} = \frac{-B_t^{br}}{T-t} dt + \sigma dW_t \quad \forall t < T, \quad B_0^{br} = 0 \quad (12)$$


where dW_t is a standard Weiner process.⁹

The restricted Brownian meander directly imposes both the equality and inequality restrictions above. We use this result and the parameter estimates from step one to simulate sequences of health status for nonhospice patients prior to hospice entry or death.¹⁰ For a patient who enters hospice prior to death, we simulate from the restricted Brownian meander with $a = \theta_0(x_i)$ and $b = \bar{\theta}(z_i, \epsilon)$; for a patient who dies without entering hospice, we simulate using $a = \theta_0(x_i)$ and $b = 0$. Figure 3 shows 50 simulated sequences for a hypothetical patient, including the continuation of the sequences during hospice enrollment.

⁹See Devroye (2010) for more details.

¹⁰We do not use nonhospice patients who neither enter hospice nor die to estimate nonhospice costs. For these patients, we have no information on the right endpoint of their health status; their health status could be almost anything. As a result, these patients are not informative about the relationship between health status and costs, so we exclude them from the nonhospice cost estimation. However, these patients are included in the estimation of the other parameters.

Figure 3: Simulated sequences of θ_t with $\sigma = 2$, $\theta_0 = 180$, $\bar{\theta} = 80$, $T_H = 200$, and $T_D = 280$.



Our second step estimator for the nonhospice cost parameters is then a simulated maximum likelihood estimator, defined as

$$(\hat{\alpha}_0, \hat{\alpha}_1, \hat{\sigma}_{nh}) = \arg \max_{\alpha_0, \alpha_1, \sigma_{nh}} \frac{1}{N} \sum_{i=1}^N \log \int_{\epsilon} \left[\frac{1}{R} \sum_{r=1}^R \prod_{t=0}^{\min\{T_{H,i}, T_{D,i}\}} \Pr(c_{i,t} | \theta_t^r, h_{i,t}) \right] \phi(\epsilon/\sigma_{\epsilon}) d\epsilon, \quad (13)$$

For the results shown below, we use $R = 25$ sequences and draw a different set of sequences for every candidate value of ϵ for every patient. We draw all of these sequences outside the estimation algorithm, which significantly reduces the computational burden.

Discussion

Our model characterizes a patient's hospice enrollment decision as trading off utility against mortality: the palliative approach of hospice care may be more comfortable and convenient for a patient nearing the end of life, but foregoing curative treatments may accelerate his death. The parameter κ determines the degree to which hospice enrollment speeds a patient's demise. This effect is inherently difficult to identify from observational data. Most patients choose hospice when they are

already very sick, so a naïve analysis that merely compares death rates of hospice vs. non-hospice patients might suggest that hospice sharply accelerates mortality. Our model addresses this challenge in two important and complementary ways. First, we model the hospice decision as being a direct function of unobservable health status. In other words, we directly model the fundamental endogeneity problem. Second, we incorporate instruments so that our estimate of hospice’s impact on mortality is driven by plausibly exogenous variation.

As explained above, the main instrument (z) we use is the overall hospice utilization rate in the patient’s county in the prior year—i.e., the fraction of hospice-eligible deaths in the county for which the decedent was enrolled in hospice at the time of death.¹¹ To understand how these instruments help identify κ , consider two patients, A and B, located in counties with high and low hospice utilization rates, respectively. Assuming the coefficient (γ) on the utilization rate is positive, our model says that A will enroll in hospice sooner than B, meaning that A will have a higher unobserved health state θ_t at the time of hospice enrollment. Suppose A enrolls in hospice N days sooner than B. If health (θ_t) declines at the same rate for hospice patients as for non-hospice patients, then A should simply be on hospice for N days longer than B before dying. If health declines faster while on hospice, then A should be on hospice longer than B, but by fewer than N days. Since health declines $(1 + \kappa)$ times more slowly for non-hospice patients than hospice patients, A’s hospice-to-death time should be longer than B’s by approximately $N/(1 + \kappa)$ days. Thus, if we observed θ_t directly, κ would be identified by the differences in hospice-to-death times for patients who enroll in hospice at different levels of θ_t due to exogenous regional variation in the popularity or availability of hospice. Obviously, we do not observe θ_t directly, so the key comparisons are observed with some noise, but the logic of the above identification argument still applies.

6 Results

6.1 Parameter estimates

Tables 6 and 7 present the results from the first and second steps of the estimation procedure described above.

¹¹These rates are calculated from hospice deaths reported in cost reports filed with CMS (the numerator) and non-accidental deaths reported by the CDC (the denominator). See Chung and Sorensen (2017) for additional details.

Table 6: First Step Estimation Results

| Parameter | Estimate | (Std. Err.) |
|--------------------------------------|----------|-------------|
| Initial health: | | |
| Age (β_1) | -2.468 | |
| HCC score (β_2) | -7.680 | |
| Hospice utility: | | |
| Constant (Δ) | 21.550 | |
| Utilization instrument (γ) | 0.0002 | |
| Standard deviation σ_ϵ | 43.274 | |
| Evolution of health: | | |
| Nonhospice drift $-(1 + \kappa)$ | -0.578 | |
| Volatility σ | 27.189 | |
| Number of patients | 29,661 | |
| Log-likelihood | -8.243 | |

Table 7: Second Step Estimation Results

| Parameter | Estimate | (Std. Err.) |
|---------------------------------------|--------------------------------------|-------------|
| Nonhospice costs: | | |
| $\lambda_{nh} = \Pr(c_t = 0 h_t = 0)$ | 0.641 | (0.0002) |
| Constant (α_0) | 8.813 | |
| $\log(\theta_t)$ (α_1) | -0.676 | |
| Standard deviation (σ_{nh}) | 2.084 | |
| Hospice costs | | |
| $\lambda_h = \Pr(c_t = 0 h_t = 1)$ | 0.110 | (0.0004) |
| Mean (α_h) | 4.932 | (0.001) |
| Standard deviation (σ_h) | 0.890 | (0.001) |
| Number of patients | 29,644 (25,019 for nonhospice costs) | |
| Number of observations | 9,474,506 | |

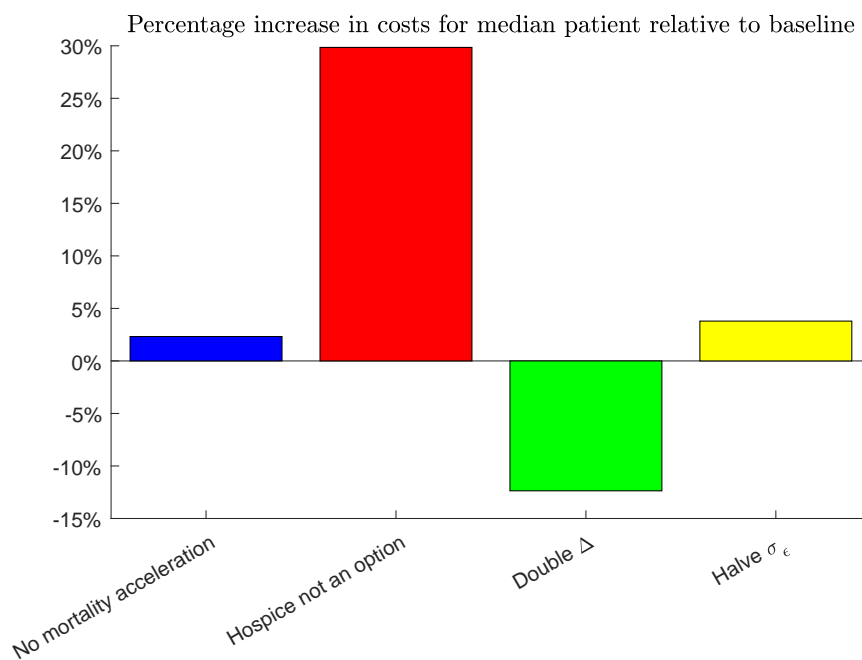
6.2 Counterfactual simulations

We generated θ_t sequences for 50,000 simulated patients and calculated their hospice enrollment decisions and simulated costs under the following scenarios:

- (1) The estimated model (benchmark case)
- (2) Hospice doesn't exist
- (3) Hospice has no mortality effect (drift rates same for hospice and non-hospice)
- (4) Double the size of Δ
- (5) Halve the size of σ_ϵ

This figure summarizes the results:

Figure 4: Counterfactual simulation results



7 Conclusions

Too early to say.

References

- Blecker, Saul, Gerard F. Anderson, Robert Herbert, Nae-Yuh Wang, and Frederick L. Brancati (2011). “Hospice care and resource utilization in Medicare beneficiaries with heart failure.” *Med Care*, 49(11):985–991.
- Chan, Tat Y., and Barton H. Hamilton (2006). “Learning, private information, and the economic evaluation of randomized experiments.” *Journal of Political Economy*, 114(6): 997-1040.
- Connor, Stephen R., Bruce Pyenson, Kathryn Fitch, Carol Spence, and Kosuke Iwasaki (2007). “Comparing hospice and nonhospice patient survival among patients who die within a three-year window.” *Journal of Pain and Symptom Management*, 33(3):238–246.
- Devroye, Luc (2010). “On Exact Simulation Algorithms for Some Distributions Related to Brownian Motion and Brownian Meanders.” In *Recent Developments in Applied Probability and Statistics*, eds. Luc Devroye, Bülent Karasözen, Michael Kohler, and Ralf Korn. Heidelberg: Physica-Verlag, pp. 1-35.
- Fang, Hanming and Edward Kung (2012). “Why do life insurance policyholders lapse? The roles of income, health, and bequest motive shocks.” NBER Working Paper #17899.
- Kelley, Amy S., Partha Deb, Qingling Du, Melissa D. Aldridge Carlson, and R. Sean Morrison (2013). “Hospice enrollment saves money for Medicare and improves care quality across a number of different lengths-of-stay.” *Health Affairs*, 32(3):552–561.
- Molini, A., P. Talkner, G.G. Katul, and A. Porporato (2011). “First Passage Time Statistics of Brownian Motion with Purely Time Dependent Drift and Diffusion.” *Physica A* 390(11):1841–1852.
- Pope, Gregory C., Randall P. Ellis, Arlene S. Ash, John Z. Ayanian, David W. Bates, Helen Burstin, Lisa I. Iezzoni, Edward Marcantonio, and Bei Wu (2000). “Diagnostic Cost Group Hierarchical Condition Category Models for Medicare Risk Adjustment.” Health Care Financing Administration report.
- Pyenson, Bruce, Stephen Connor, Kathryn Fitch, and Barry Kinzbrunner (2004). “Medicare cost in matched hospice and non-hospice cohorts.” *Journal of Pain and Symptom Management*, 28(3):200–210.
- Taylor Jr., Donald H., Jan Ostermann, Courtney H. Van Houtven, James A. Tulsky, and Karen Steinhauser (2007). “What length of hospice use maximizes reduction in medical expenditures near death in the US Medicare program?” *Social Science & Medicine*, 65(7):1466–1478.