

WHY IS END-OF-LIFE SPENDING SO HIGH? EVIDENCE FROM CANCER PATIENTS

Dan Zeltzer, Liran Einav, Amy Finkelstein, Tzvi Shir, Salomon M. Stemmer, and Ran D. Balicer*

Abstract—We study the sources of high end-of-life spending for cancer patients. Even among patients with similar initial prognoses, spending in the year postdiagnosis is over twice as high for those who die within the year than those who survive. Elevated spending on decedents is predominantly driven by higher inpatient spending, particularly low-intensity admissions. However, most such admissions do not result in death, making it difficult to target spending reductions. Furthermore, end-of-life spending is substantially more elevated for younger patients, compared to older patients with similar prognoses. Results highlight sources of high end-of-life spending without revealing any natural “remedies.”

I. Introduction

MEDICAL spending is highly concentrated at the end of life. A widely cited fact is that, in the United States, only 5% of Medicare beneficiaries die each year, but one-quarter of Medicare spending occurs in the last 12 months of life (Riley & Lubitz, 2010). This is frequently touted as indicative of obvious waste and inefficiency: we spend a large share of healthcare dollars on individuals certain to die within a short period (e.g., Emanuel & Emanuel, 1994; Medicare Payment Advisory Commission, 1999).

In this paper, we ask: why is spending concentrated at the end of life? Our approach is motivated by existing work that has already ruled out two natural hypotheses. One is that high end-of-life spending reflects idiosyncratic inefficiencies embodied in the specific institutional features of the US healthcare system. This is not the case. Healthcare spending is similarly—or more—concentrated at the end of life in other OECD countries (French et al., 2017). Another is that the focus on high end-of-life spending is misguided due to classic hindsight bias (Fischhoff, 1975): we spend more on the sick, and the sick are more likely to die, which together accounts for the concentration of spending on those who die. While this qualitative statement is (naturally) true, it cannot explain the quantitative patterns: even conditioning on initial health, spending on decedents is still over twice as high as that on survivors (Einav et al., 2018).

Received for publication January 10, 2020. Revision accepted for publication April 26, 2021. Editor: Benjamin R. Handel.

*Zeltzer: Tel Aviv University; Einav: Stanford University and NBER; Finkelstein: MIT and NBER; Shir: Clalit Health Services; Stemmer: Rabin Medical Center and Tel Aviv University; Balicer: Clalit Health Services and Ben Gurion University.

We thank seminar and conference participants at Princeton University, University of Pennsylvania, the American Society of Health Economists meeting, and the World Congress of the Econometric Society. We acknowledge financial support from the Koret Foundation. Dan Zeltzer acknowledges financial support from the Pinhas Sapir Center for Development. Avichai Chasid, Michael Leshchinsky, and Joseph Rashba provided excellent research assistance.

A supplemental appendix is available online at https://doi.org/10.1162/rest_a_01066.

To shed light on the sources of elevated spending on decedents compared to *ex ante* similar individuals who survive, we focus our analysis on a specific set of individuals: patients newly diagnosed with cancer. Focusing on a specific disease provides us with a relatively more homogeneous set of conditions and treatment options, thereby allowing us to dig deeper into the nature of spending on decedents compared to survivors, albeit on a subset of the population. Patterns of end-of-life spending for cancer patients are broadly similar to those in the general population: spending is elevated at the end of life across a range of OECD countries (Bekelman et al., 2016; French et al., 2017) and, as we will show, this elevated end-of-life spending occurs even across patients with the same initial mortality prognosis when the cancer is detected.

Cancer is a particularly useful disease to focus on for several reasons. First, it is common and expensive. Cancer is the second-leading cause of death in developed countries—accounting for over one-fifth of deaths—and treatment options are resource-intensive (Emanuel et al., 2002; Heron, 2013; Bekelman et al., 2016). Second, cancer has a clear diagnosis date, after which major spending decisions occur over a relatively short period. This makes it easier to analyze the course of spending on cancer than on other diseases, such as hypertension, for which the diagnosis date and treatment period are less clearly defined. Third, the treatment options for cancer can be classified into a few, discrete treatment options which patients may move between (e.g., surgery, outpatient chemotherapy, radiotherapy, maintenance care); this allows us to examine how treatment decisions change as the mortality prognosis evolves. Fourth, cancer unfortunately affects a wide age range, which allows us to compare treatment patterns between younger and older individuals who have very different residual life expectancy conditional on successful treatment.

We analyze detailed and comprehensive longitudinal medical data covering about half of the Israeli population from 2000–2016. The data come from Clalit Health Services, the largest of four HMOs in Israel that provide universal, tax-funded health insurance to all residents. The data include electronic medical records (EMR) as well as claims data. They therefore permit a richer set of health measures than are available in the US Medicare claims data, in which end-of-life spending has been extensively analyzed (Barnato et al., 2004; Nicholas et al., 2011; Morden et al., 2012; Teno et al., 2013; Einav et al., 2018). In addition, the data allow us to analyze end-of-life spending patterns over the entire age range of patients, rather than limiting ourselves to the elderly.

Our primary focus is on 160,000 adults (ages 25 and older) who were newly diagnosed with cancer from 2001 through

2013. These cancer patients have a 20% annual mortality rate, much higher than the 1.2% annual mortality rate in our overall adult population. For each individual in the data, we generate a prediction of the probability that they will die in the year following their diagnosis; we refer to this as the patient's "initial prognosis." We also generate a separate mortality prediction following each major clinical event (such as a hospital admission or an outpatient chemotherapy spell) during the course of their treatment in the first year postdiagnosis; we refer to these predictions as the patient's "current prognoses." To generate these prognoses, we apply standard machine learning techniques to a rich dataset with hundreds of potential predictors, including demographics, healthcare utilization, diagnoses, and various biomarkers in the prior 12 months. We analyze average monthly spending and healthcare use in the 12 months post-cancer diagnosis (or post-major clinical event) for ex-post-survivors (i.e., those who remain alive 12 months after their cancer diagnosis) compared to ex-post-decedents (those who die within 12 months of their cancer diagnosis), limiting attention to months in which decedents (and likewise survivors) are alive.

We have three main findings that together provide insight into the sources of elevated end-of-life spending. First, elevated spending on decedents relative to survivors with the same initial prognosis is almost entirely driven by elevated inpatient spending, particularly low-intensity admissions with few procedures. Although inpatient spending is only 40% of medical spending among survivors, higher spending on inpatient care accounts for 95% of the elevated spending on decedents. Spending on all other care—including outpatient care, radiation, and chemotherapy—is only slightly larger among decedents. Within inpatient care, spending on low-intensity admissions accounts for only one quarter of inpatient spending among survivors, but for about two thirds of the elevated inpatient spending on decedents.

Second, treatment patterns are consistent with a switch to maintenance inpatient care at the end of life. In particular, a sharp worsening of the current prognosis is associated with an increase in low-intensity admissions. As a result, for decedents, spending on low-intensity admissions tends to spike in what is (ex post) the last few months of life, regardless of survival duration, while spending on chemotherapy and radiation tends to spike right after the initial diagnosis and tails off in the last few months, again regardless of survival duration. Nonetheless, a large share of low-intensity admissions do not end in death within the subsequent two months—even among patients with a poor prognosis at the time of admission—suggesting that it is not easy to *ex ante* identify what ex post is spending at the end of life.

Third, we find that among patients with the same initial prognosis, the elevated spending for decedents is particularly pronounced for younger patients. This pattern also holds within cancer type (so that we are comparing across patients for whom the available "technology" or treatment options are broadly similar). Since a key difference across patients

of different ages is life expectancy conditional on successful cancer treatment, these age patterns suggest that treatment decisions may not simply reflect a short-run goal of staving off near-term mortality, but may be affected by considerations that take into account a longer-run horizon.

The rest of the paper proceeds as follows. Section II presents a brief conceptual framework designed to clarify what we are able to measure relative to the fundamental objects of interest. Section III describes our setting, data, and the construction and performance of our initial prognosis algorithm. Section IV presents the results. The last section concludes.

II. Conceptual Framework

As with most work in (health) economics, we do not measure the direct objects of interest. We therefore briefly clarify what those fundamental objects are, and how the objects of our analysis relate to them.

Consider a population of individuals, each denoted by i . Absent any spending (i.e., treatment), individual i is associated with a baseline death probability of θ_i , which is drawn from a distribution $G(\theta_i)$. Let $f(\theta, s)$ define the health production function, which maps individual baseline death probability (pretreatment) to what the death probability would be when medical spending is s . By definition, $\theta_i \equiv f(\theta_i, 0)$. It is also natural to assume that spending is (weakly) productive for all individuals (i.e., $\frac{\partial f}{\partial s} \leq 0 \forall \theta$), and that at any level of spending s , the order of risk across types is preserved (i.e., $\frac{\partial f}{\partial \theta} > 0 \forall s$).

If we were able to measure the mortality risk in the absence of treatment $G(\theta_i)$ and the health production function $f(\theta, s)$, we would be able to determine the optimal spending policy $s(\theta)$. To see this, consider for example a social objective to minimize overall mortality, subject to a budget constraint B .¹ The social planner's problem is thus to solve:

$$\min_{s(\theta)} \int f(\theta, s) dG(\theta) \quad s.t. \quad \int s(\theta) dG(\theta) \leq B. \quad (1)$$

However, of course, these two key objects are inherently difficult to observe. The mortality risk in the absence of treatment $G(\theta_i)$ —the so-called "natural history" of the disease—is almost never observed, because the sick almost always receive treatment. The health production function $f(\theta, s)$ is arguably the most sought-after object in health economics, yet empirical knowledge of it is sorely lacking.

Since we cannot observe the objects of interest, we instead construct estimates of what we can observe. Specifically, instead of the mortality risk in the absence of treatment $G(\theta_i)$,

¹This is of course merely an illustrative example. Naturally, one could consider alternative social objectives, such as assigning different weights to individuals by age.

FIGURE 1.—ILLUSTRATION OF THE CONCEPTUAL FRAMEWORK

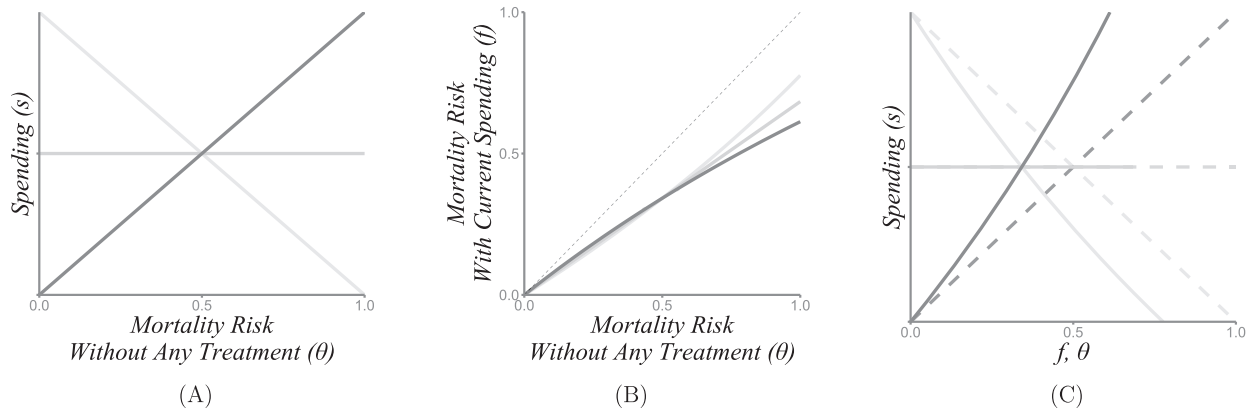


Figure shows an illustration of the conceptual framework discussed in section II. The panels show statistics for three different spending policies. The policies, marked by different colors, are: spending concentrated on the healthy (light gray), spending concentrated on the sick (dark gray), and uniform spending (gray). Panel A shows spending, s , as a function of (unobserved) individual type θ , which is defined by death probability in the absence of treatment. Panel B shows actual mortality with treatment, $f(\theta, s(\theta))$, as a function of type θ . The dashed line in this panel is the identity (45°) line. Panel C shows spending as a function of observed actual mortality with treatment, f , in solid lines, along with the underlying policies from panel A, which are unobserved, repeated in dashed lines. Spending is normalized to have a 0–1 range and has no units. See section II for details of the calculations.

we measure the equilibrium distribution of mortality risk, $H(f(\theta, s(\theta)))$. This, of course, is endogenous to the healthcare spending policy $s(\theta)$. Instead of measuring this spending policy $s(\theta)$, we likewise measure the relationship between spending and equilibrium mortality risk, $s(f(\theta, s(\theta)))$.

Under some assumptions, this (endogenous) object $s(f(\theta, s(\theta)))$ can still be informative of the deeper economic primitive of interest $s(\theta)$. For example, imagine that $G(\theta)$ is a uniform distribution over $[0, 1]$, and that $f(\theta, s) = 1 - s^{\frac{1}{2}}\theta$ (where $s \in [0, 1]$), so that the health returns to spending are increasing in baseline mortality risk.²

Figure 1 illustrates how, under these assumptions, the objects of interest relate to the ones we will measure. In panel A, we consider three possible (budget-neutral) shapes to the healthcare spending function $s(\theta)$: uniform spending on all types, a spending policy that favors the sick, and a spending policy that favors the healthy. Panel B shows the implications of these different policies for the way health improves differentially by θ . Given our assumption about the health production function, the figure makes clear that the optimal spending policy would be for spending to be increasing in mortality risk θ .

Panel C of figure 1 presents the implications of these different spending policies for the relationship between spending and post-treatment mortality risk, $f(\theta, s(\theta))$. It is these last objects that are estimable and the focus of our empirical analysis. The figure illustrates that the (endogenous) object $s(f(\theta, s(\theta)))$ can still be informative about the deeper economic primitives and, in particular, about the health spending policy $s(\theta)$ (at least under strong assumptions restricting the types of spending policies we consider).

²Again, this is merely an illustrative example. One could of course assume a different health production function, such as one in which there were higher returns to spending for lower-mortality individuals.

III. Data and Methods

A. Setting and Data

Our data come from Clalit Health Services, the largest of Israel's four nonprofit Health Maintenance Organizations (HMOs) that provide universal tax-funded healthcare coverage from birth to all Israeli residents, in accordance with the National Health Insurance Law (1995). Under Israeli health insurance, covered services are essentially fully subsidized by risk-adjusted capitated payments from the government.³ The coverage broadly resembles that of U.S. Medicare Parts A, B, and D, and includes hospital admissions, outpatient services, physician consults, drugs, and durable medical equipment.

Clalit Health Services is an integrated provider and insurer. It provides most of the services it finances and reimburses preauthorized services purchased from external providers. Its members are admitted to all of Israel's thirty general hospitals, eight of which Clalit directly owns and operates. It employs over 11,000 physicians and 10,000 nurses, operates over 1,500 primary clinics across the country, and provides multiple outpatient services. By 2001, Clalit had adopted electronic medical records (EMRs) for its enrollees.

The data cover a large and stable population. Clalit covers about half of the Israeli population, approximately 4.5 million members of all ages. Churn is extremely low: each year, less than 1% of Clalit enrollees switch to another HMO. Thus, most adults remain enrolled with Clalit throughout their lifetime. Appendix A provides more detail on the Israeli Health Insurance System and on our particular data provider, the insurer Clalit.

The data are available longitudinally (from 2000 through 2016) across all possible care settings. They are rich and

³There are no premiums, small copays for outpatient services and emergency room visits, no copays for admissions, and a maximum out-of-pocket cap of 800 New Israeli Shekels (NIS, or about USD 200) per quarter.

detailed. As with the US Medicare data, they contain claim-level data on patient encounters, diagnoses and payments, demographics, and date of death if any. In addition, through the EMR, we also observe a rich set of lab results, screening, imaging, and health measures that are not available in standard claims data, including, for example, vital signs, blood tests, and body mass index.

We supplement these data with linked data on the exact timing of the first diagnosis of cancer from the Israel National Cancer Registry, to which reporting has been mandatory since 1982. While this information can also be extracted from claims data, the Registry provides an official diagnosis date. We also take advantage of EMR data from admissions, for the subset of admissions in Clalit-owned hospitals for which such data are complete, to characterize the types of procedures performed for different admission categories; Clalit-owned hospitals comprise about 40% of admissions.

B. Sample and Key Variables

We focus our analysis on adults (25 and older) who had a new cancer diagnosis between 2001 and 2013. We restrict the sample to patients with at least one year of coverage prior to their initial diagnosis and who remain with Clalit for at least 12 months after the diagnosis date (or until death); these restrictions exclude less than 1% of patients. For the small fraction of patients who are associated with multiple (distinct) cancer diagnoses during the observation period, we restrict attention to the first diagnosis. For comparative purposes, we also present some descriptive statistics for the full population of all 2.3 million adults (25 and older) covered by Clalit as of January 1, 2013, with the (minor) sample restriction that they are observed for at least one year prior to and one year subsequent to that date (or until death).

Outcomes. We focus primarily on one-year mortality and average monthly healthcare spending and healthcare use over this one year. Spending measures are obtained from the administrative records of Clalit. We observe payments for all services detailed in encounter-level claims data (including inpatient admissions, emergency department visits, treatments and diagnostic services provided in outpatient clinics, both within and outside hospitals, and prescription drug purchases).⁴

Our main spending measure is adjusted average monthly spending. This measure averages spending only over months in which the patient is alive, in order to account for the shorter survival duration of decedents. This is useful when com-

paring spending patterns between decedents and survivors. Specifically, adjusted average monthly spending is defined as

$$\bar{y}^I = \frac{\sum_{i \in I} y_i}{\sum_{i \in I} (T_i/30)}, \quad (2)$$

where I is a set of individuals, y_i is total healthcare spending of individual i in the 12 months following the index date, and $T_i \in (0, 365]$ is the right-censored number of days individual i survived after the index date.

We also construct several measures of the nature of inpatient admissions. We classify all admissions based on whether they are unplanned (i.e., originated through the emergency room) or planned. We also classify them as high or low “intensity,” with high versus low intensity defined based on the average daily spending for different hospital wards (i.e., hospital units). As would be expected, the high-intensity wards, such as general surgery, tend to have a much higher share of admissions with surgical procedures than low-intensity wards, such as oncology or internal medicine (appendix table A1). Finally, for the 40% of admissions in which we can observe inpatient procedures, we measure whether the admission involved each of six (nonmutually exclusive) different types of inpatient procedures: diagnostics (lab and imaging), surgeries, inpatient chemotherapies, inpatient radiation therapies, maintenance (e.g., evaluation, feeding, pain management), and all others.⁵

Mortality predictors. We exploit the richness of the data to code hundreds of potential mortality predictors that we use as features that go into training our prognosis algorithms; appendix B.1 describes these predictors and their construction in detail. Broadly speaking, they fall into four main categories. First, we use demographic data from administratively sourced information on birth date, gender, social security transfers, disability, and location-based socioeconomic status. Second, we measure monthly healthcare utilization and spending by type of service in the claims data. Third, we calculate measures of overall morbidity based on all diagnoses documented in clinical encounters. Specifically, we use the Johns Hopkins Adjusted Clinical Groups (ACG) system to predict resource utilization and the probability of major health events.⁶ All of these measures are standard in claims data.

Our fourth category of variables is less commonly available: the EMR data provide additional health measures. These include BMI, vital signs measures, blood test results, and

⁴The spending measures represent actual payments made by Clalit, not list charges. Even in cases where the hospital is owned by Clalit, it serves as a separate financial entity as Clalit hospitals also serve non-Clalit patients and charge other insurers similar prices. We do not directly observe spending for office-based consults provided by salaried physicians in Clalit-owned clinics. For these visits, we construct per-visit charges that are based on customary charges by nonemployed providers; these comprise about 2.8% of total spending in our cancer sample.

⁵As described earlier, we only observe inpatient procedure data for patients admitted to Clalit-owned hospitals. The characteristics of patients admitted to Clalit-owned hospitals are similar to those admitted to other hospitals (appendix table A2).

⁶This system is used by both commercial insurers and noncommercial healthcare organizations worldwide (as well as by Clalit) to describe or predict a population’s past or future healthcare utilization and costs. For more information, see *The Johns Hopkins ACG System Version 11.0 Technical Reference Guide* (2014).

TABLE 1.—DEMOGRAPHICS, COST, AND MORTALITY

	Cancer sample			General population sample		
	All (1)	Decedent (2)	Survivor (3)	All (4)	Decedent (5)	Survivor (6)
Characteristics						
Age (mean)	65	73	64	50	78	50
Female (%)	52.1	44.8	53.9	52.4	52.1	52.4
High SES zip-code (%)	23.4	18.7	24.5	21.4	18.8	21.5
Supplementary insurance ^a (%)	70.1	54.7	73.9	74.8	59.8	75.0
Mortality rate						
1 month (%)	3.7	19.0	—	0.1	10.1	—
1 year (%)	19.5	100.0	0.0	1.2	100.0	0.0
3 years (%)	32.4	—	16.0	3.5	—	2.4
Utilization						
12 months before index date						
Average monthly spending (NIS)	1,406	2,290	1,192	484	3,648	446
Any admission (%)	53.3	75.7	47.9	12.1	55.0	11.6
12 months after index date						
Average monthly spending (unadjusted NIS)	4,730	4,987	4,668	556	4,178	514
Average monthly spending (adjusted NIS) ^b	5,380	13,139	4,668	560	8,638	514
Any admission (%)	73.0	87.8	69.5	12.8	78.8	12.0
Number of beneficiaries	166,839	32,518	134,321	2,372,582	27,673	2,344,909

Table shows descriptive statistics for our main sample—the adult population diagnosed with cancer (columns 1–3), and for comparison, the overall adult population (columns 4–6). This table (and appendix figure A4) describes the full sample, which we later split into training and test sets. All other exhibits are based on the test set. Columns 1 and 4 show statistics for all patients; columns 2 and 5 show statistics for ex-post decedents, that is, those who died within 12 months after the index date; columns 3 and 6 show statistics for ex-post survivors, that is, those who remain alive after 12 months. The index event is defined as the date of initial diagnosis for cancer patients, and January 1, 2018 for the general population. By definition, the mortality rate within one year of the initial prognosis is 100 for decedents and 0 for survivors. Utilization measures are shown for the periods of 12 months before and 12 months after the index date. All spending measures are in current New Israeli Shekels (NIS); during our study period the exchange rate was about 4 NIS per USD. ^aSee appendix A for more details. ^bAdjustment is based on equation (2). See text for more details.

information on drug adherence. We also measure the cancer topography (i.e., body part type) from the national cancer registry data.

We use these predictors to form two types of mortality predictions. First, for each patient, we predict one-year mortality risk at the date of diagnosis; we refer to this as the “*initial prognosis*.” For this measure, all of the healthcare and health measures in the EMR and the claims data are measured on or up to 12 months prior to the diagnosis date.

Second, we generate one-year predicted mortality risk at the start of each of five major clinical events (which cover the major broad categories of cancer care): high-intensity hospital admission, low-intensity hospital admission, emergency room visit (which may mark an unexpected deterioration), outpatient drug therapy spell, or outpatient radiation therapy spell. We refer to the one-year predicted mortality rate at the start of a clinical event as the “*current prognosis*.” For this measure, all of the predictors are measured on or up to the 12 months prior to the clinical event. The predictors include the diagnosis (cancer topography) itself as well as the sequence of major clinical events postdiagnosis.

Summary statistics. Table 1 presents summary statistics for the 160,000 cancer patients, with statistics for the general population also shown for comparison. Cancer patients are on average older and sicker than the general population, even before they get diagnosed with cancer. The one-year mortality rate for cancer patients (19.5%) is much higher than that of the general population (1.2%); one fifth of cancer decedents (who die in the year following their diagnosis) die within a month of diagnosis. Those cancer patients who survive a year have a much lower mortality rate in subsequent years; only

81% of cancer patients survive a full year, but 84% of those survive an additional two years.

We also compare decedents (who die within a year of diagnosis) to (one-year) survivors. Decedents are sicker and more expensive than survivors, even before a cancer diagnosis (table 1, bottom panel). They have more hospital admissions and spend on average more than survivors in the 12 months prior to diagnosis. In the year leading to a cancer diagnosis, decedents spend on average NIS 2,300 (approximately USD 575) per month; survivors spend NIS 1,200 (approximately USD 300) per month. Decedents are also older than survivors (73 versus 64 years old on average).⁷

For some of our analyses, we analyze how prognoses and spending decisions change over the course of treatment. To do so, we limit our sample of patient-events to those that have at least one clinical event following the initial cancer diagnosis and analyze outcomes at the event level. Patients remain in this sample until death or remission and will show up multiple times if they have more than one clinical event following the initial diagnosis. This allows us to focus on the subset of patients who remain in treatment and therefore require further medical decision making. The resulting sample has a total of 292,484 patient-event observations, with 2,610 distinct sequences of between one and seven clinical events.⁸

⁷Appendix table A3 shows statistics further disaggregated by type of cancer. Breast, prostate, and colon cancer are the three most common cancers, collectively accounting for about one-third of all cancer diagnoses. Mortality rates and spending vary substantially across types of cancer. While we pool all cancer types to generate our main results, cancer type is always included in our mortality prediction algorithm. We will report below on some analyses that are performed separately by cancer type.

⁸For expositional clarity, we include only the first seven events for each patient. Less than 2% of patients have additional events.

FIGURE 2.—DISTRIBUTION OF NUMBER OF CASES BY TYPE AND NUMBER OF MAJOR CLINICAL EVENTS

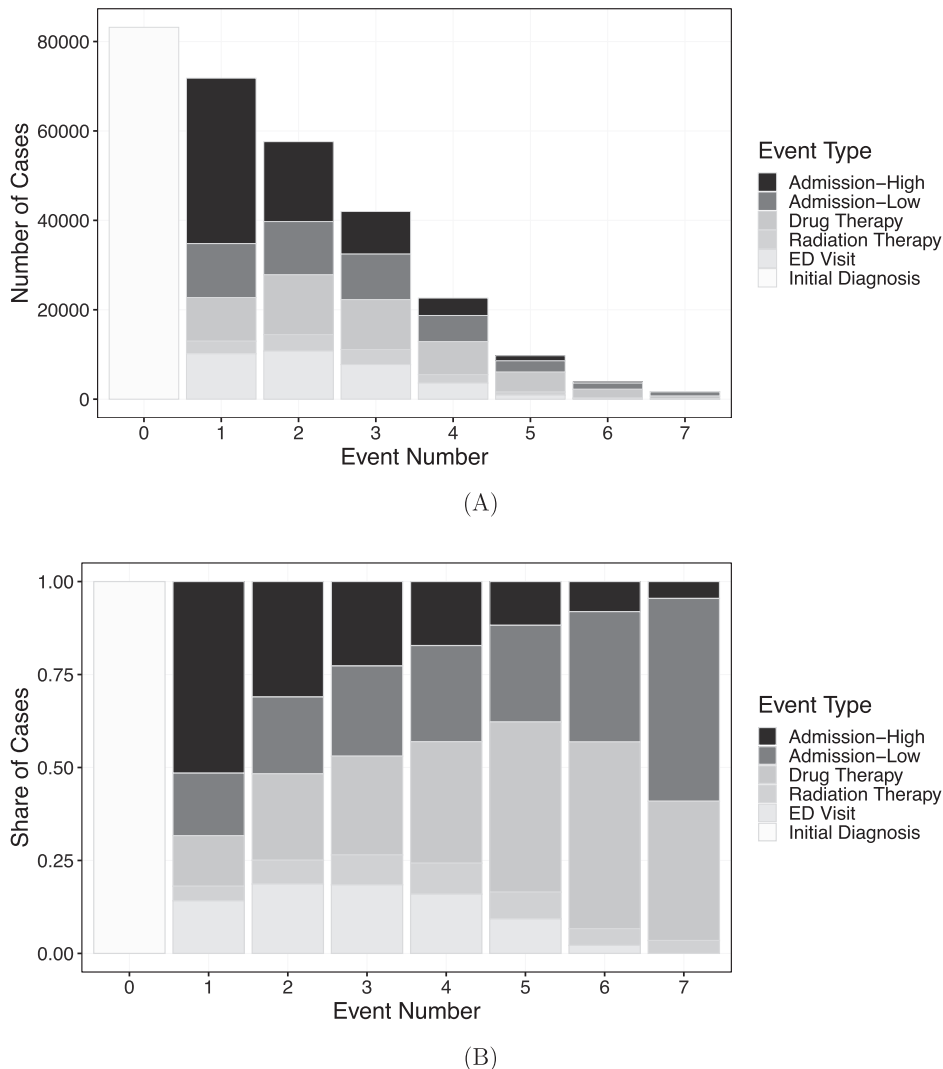


Figure shows, for the sample of major clinical events of cancer cases ($N = 292,484$ patient-events), the number (panel A) and share (panel B) of cases still in treatment and their most recent major event, as a function of the sequential number of this event. Colors denote the type of the most recent event. Admission-High and Admission-Low denote high- and low-intensity admissions. Drug Therapy is a spell of either chemotherapy or biological drug treatment. Radiation Therapy denotes a spell of such therapy. ED visit is emergency department visit that did not result in an admission to a hospital. Initial Diagnosis denotes initial cancer diagnosis.

Figure 2 shows the distribution of event types among all cases still in treatment, after different (sequential) numbers of major clinical events (0 is the initial diagnosis, for all cancer patients; 1 is the first event, for all patients who had one or more events; 2 is the second event, for all patients with at least two events, etc.). Half of all cancer patients in our sample had at least three major clinical events during the year after diagnosis; a quarter of patients had at least four (panel A). High-intensity admissions (e.g., for surgical excision of solid tumors) account for more than half of the first clinical events, and more than a quarter of the second clinical events (panel B). This share declines for subsequent events, giving way to an increasing share of outpatient drug therapies and low-intensity admissions. This increase in the share of low-intensity admissions is concentrated among patients with the deadlier cancer types (brain, lung, and pancreas; see appendix figure A1).

C. Prognosis Algorithms

We apply standard machine learning techniques to the rich dataset with hundreds of potential predictors described in the preceding section to create our one-year mortality predictions (both “initial prognosis” and “current prognosis”). To model and estimate mortality risk, we use Extreme Gradient Boosting (Chen & Guestrin, 2016), a popular sequential ensemble method that iteratively and greedily constructs a series of classifiers, with each classifier being used to fit the residuals of the previous classifier. This method can flexibly accommodate interactions among predictors and fit an arbitrary differentiable criterion function.

To avoid over-fitting, we follow standard practice and randomly split our original sample into two equally sized samples: the “test sample,” which we do not use as we optimize our prediction algorithm, and the “training sample,” which we use to fit our predictive model. The training sample is

used only for fitting the predictive model. We tune key parameters by fivefold cross-validation to maximize the area under the curve (AUC) criterion. The trained model is then used to predict mortality in the testing sample, over which the rest of the analysis is performed. Unless otherwise noted, all exhibits are based on the test sample. Appendix B.2 provides more detail on the construction and performance of the algorithms.

The prognoses generated by the algorithms are the empirical analog of the equilibrium distribution of mortality risk $H(f(\theta, s(\theta)))$ in section II. In what follows, we graphically analyze spending patterns as a function of these prognoses (i.e., $s(f(\theta, s(\theta)))$).

In addition, we also compare spending patterns for ex-post decedents and ex-post survivors with the same prognosis. To quantify outcome differences for survivors and decedents with the same prognosis, we report differences in outcomes between decedents and a reweighted distribution of survivors, reweighted so that they have the same distribution of prognoses as decedents. Namely,

$$\bar{y}_{\text{survivor(reweighted)}} = \int y_{\text{survivor}}(\mu) d\mu_{\text{decedent}}, \quad (3)$$

where y_{survivor} denotes monthly survivor spending, and μ_{decedent} is a measure of decedent risk. In our baseline analysis, μ_{decedent} is a two-dimensional distribution of prognosis and months since diagnosis.⁹ We reweight spending by months since diagnosis in addition to prognosis because both spending and mortality tend to be concentrated early in the year postdiagnosis. Appendix figures A2 and A3 show the raw data underlying the reweighting procedure.

IV. Results

A. Patterns of End-of-Life Spending and Mortality Risk

Healthcare spending for cancer patients is disproportionately concentrated on decedents. For example, the share of spending on decedents relative to survivors is almost three times higher than decedents' share of days lived (appendix figure A4).¹⁰ However, the *ex ante* differences between decedents and survivors shown in table 1 highlight the need to adjust for mortality prognosis when discussing spending differences between decedents and survivor.

These mortality prognoses show that it is very hard to predict who will die within the coming year (which is consistent

with similar findings for a general population [Einav et al., 2018]). For example, the 95th percentile of the initial prognosis for cancer patients is an annual mortality rate of only 81%, and only one quarter of those who end up dying within the year have initial mortality prognoses greater than 80%. Moreover, individuals with very poor initial prognoses account for only a very small share of total spending (appendix figure A5). For example, less than 10% of spending on cancer patients is accounted for by individuals with initial predicted mortality above 80%. Even among pancreatic cancer patients, who have the highest annual mortality rate (two thirds), less than 5% of patients have an initial annual mortality prognosis above 95%, and less than 55% of those who end up dying within the year have initial mortality prognoses greater than 80% (appendix table A3). These findings underscore a fundamental point: there is no sizable mass of cancer patients for whom, at the time of initial diagnosis, death is certain or "near certain" (within the year).¹¹

An obvious explanation for the concentration of spending at the end of life is that spending is higher among sicker patients, and sicker patients are also more likely to die. We therefore examine spending patterns by initial prognosis (figure 3). In the year following diagnosis, unadjusted average monthly spending—which includes month after death when spending is mechanically zero—shows an inverted U-shaped pattern with respect to initial prognosis. This is driven by the fact that higher mortality-risk individuals survive on average for fewer months. This is why in the remainder of the paper we focus on adjusted average monthly spending (which averages only over months alive). Adjusted average monthly spending is strongly increasing in initial mortality risk, presumably reflecting the fact that spending is higher for sicker patients. However, even after conditioning on initial prognosis, adjusted average monthly spending is elevated for decedents compared to survivors (panel B). This elevation of spending on decedents relative to survivors with the same initial prognosis is particularly pronounced for patients with good initial prognoses (i.e., low predicted mortality).

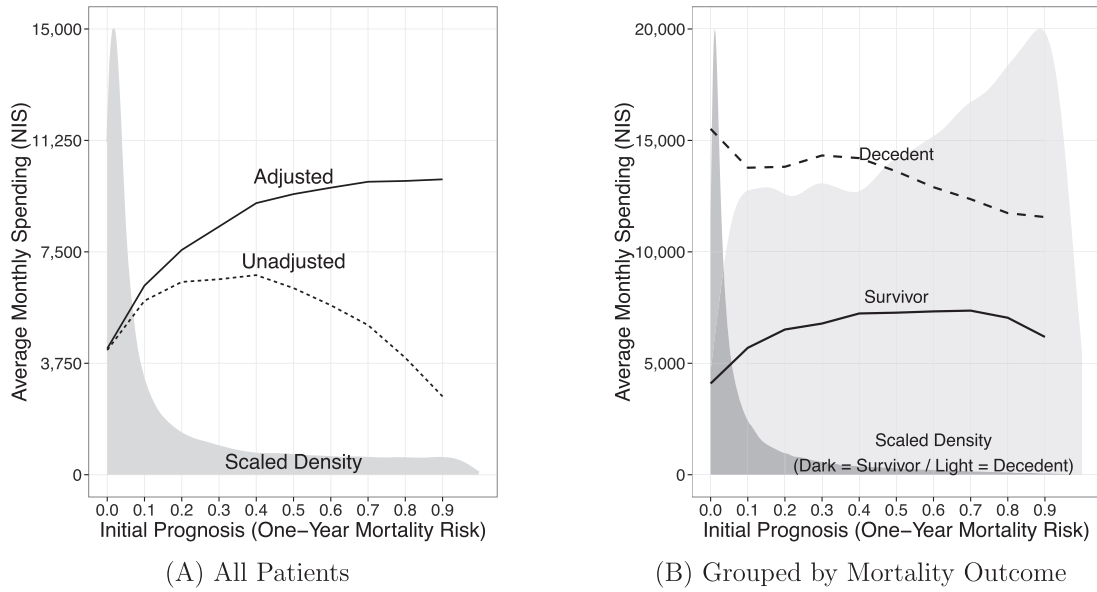
The first row of table 2 quantifies the difference in spending between decedents and survivors. Without adjusting for differences in the initial prognosis, decedents' average adjusted monthly spending is nearly three times greater than survivors' (NIS 13,204 versus 4,671). Reweighting survivor spending by decedent risk at the time of diagnosis (column 2), the gross difference of NIS 8,533 drops to 5,372. In other words, differences in initial prognosis between ex-post decedents and ex-post survivors account for almost two-fifths of

⁹We approximate this integral by partitioning prognoses into ten equally sized bins and partitioning the year to 12 months. We calculate the mean survivor spending in each mortality probability-month bin. We then average across all bins, using the number of decedents in that prognosis-month bin as weights. Note that, by construction, $\bar{s}_{\text{decedent}} = \int s_{\text{decedent}}(\mu) d\mu_{\text{decedent}}$, so we only reweight survivor spending.

¹⁰Not surprisingly, spending on decedents is somewhat less concentrated in the cancer population than for the general adult population (for whom, appendix figure A4 shows the decedent share of spending is fourteen times higher than their share of days lived). This is because virtually all cancer patients receive some nontrivial amount of medical care, while many adults in the general population receive no care.

¹¹This fact is true even when we restrict attention to the subsample of 23,000 patients whose entire hospital care was furnished in Clalit-owned hospitals, for which we have full EMR coverage. For example, the 95th percentile of the initial prognosis (predicted one-year mortality risk) is 81.1 in the full study sample of cancer patients and 81.5 among cancer patients whose care is concentrated in Clalit-owned hospital. For pancreatic cancer patients who end up dying within the year, 51.4% have initial mortality prognoses greater than 80%; for those patients whose care is concentrated in Clalit-owned hospital, the corresponding share is 54.0%.

FIGURE 3.—SPENDING BY PREDICTED MORTALITY



Figures show the distribution of initial prognosis (one-year mortality risk) and average adjusted monthly spending in the 12 months post-initial cancer diagnosis as a function of initial prognosis. $N = 83,181$ patients. Scaled density (in gray) is the kernel density estimate of the probability density function of the mortality prognosis (which integrates to one), scaled to fit the plot height. Panel A shows data for all patients combined. Unadjusted spending (dashed line) is average monthly spending, calculated over the entire year following a cancer diagnosis, including months after death with zero spending. Adjusted spending (solid line) is the average spending over the period each patient was alive during the first year after the cancer diagnosis (see equation [2]). Panel B shows adjusted average monthly spending, separately for Survivors (solid line), defined as those patients who survived for at least one year from the index date and Decedents (dashed line), defined as those who did not. Decedent spending is adjusted for survival duration (see equation [2]). The shaded areas show scaled densities of predicted mortality for each of these groups (in light gray for decedents and in dark gray for survivors). All spending measures are in current New Israeli Shekels (NIS).

TABLE 2.—AVERAGE MONTHLY SPENDING

Category	Survivor		Decedent (3)	Difference	
	Unweighted (1)	Reweighted by decedent risk (2)		Decedent-survivor (reweighted) (4)	Percent of total difference (5)
Total	4,671	7,833	13,204	5,372	100.0
All inpatient:	1,735	4,070	9,152	5,083	94.6
Planned	1,326	2,904	5,133	2,229	41.5
Unplanned	409	1,166	4,019	2,854	53.1
Low intensity	482	1,738	5,302	3,564	66.3
High intensity	1,254	2,332	3,850	1,518	28.3
Other services:	2,936	3,763	4,052	289	5.4
Drugs	1,119	1,562	1,733	172	3.2
Outpatient	1,239	1,491	1,566	75	1.4
Imaging	191	253	222	-31	-0.6
Other	387	456	530	74	1.4

Table shows average monthly spending in the 12 months post-cancer diagnosis. Columns show results separately for decedents and survivors. Decedent spending is adjusted for survival duration (see equation [2]). Survivor spending in column 2 is reweighted by decedent risk and month-from-diagnosis (see equation [3]). Decedent-survivor is the difference between Decedent and Survivor (reweighted) spending. All spending measures are in current New Israeli Shekels (NIS). First row shows total healthcare spending, and subsequent rows show various partitions. All inpatient refers to spending on all services that are delivered during hospital admissions, and Other Services refers to spending on all services that are not part of an admission. Within inpatient, we partition into low intensity versus high intensity, and unplanned versus planned. Low intensity refers to admissions into one of four wards: Internal Medicine, Oncology, Rehabilitation, and Geriatric, which appendix table A1 shows involve the lowest average daily admission and few surgeries; High intensity is admission to all other wards. Unplanned refers to admissions through the emergency department; Planned refers to all other admissions. Within Other Services we partition into Outpatient, Drugs, Imaging, and Other. Outpatient, Drugs, and Imaging refer to hospital outpatient services, prescription drugs (except those administered during admissions), and diagnostic radiology services not during an admission, respectively. $N = 83,181$ patients.

the elevated spending on decedents. The next section explores the sources of this higher spending.

B. Sources of Elevated Spending on Decedents

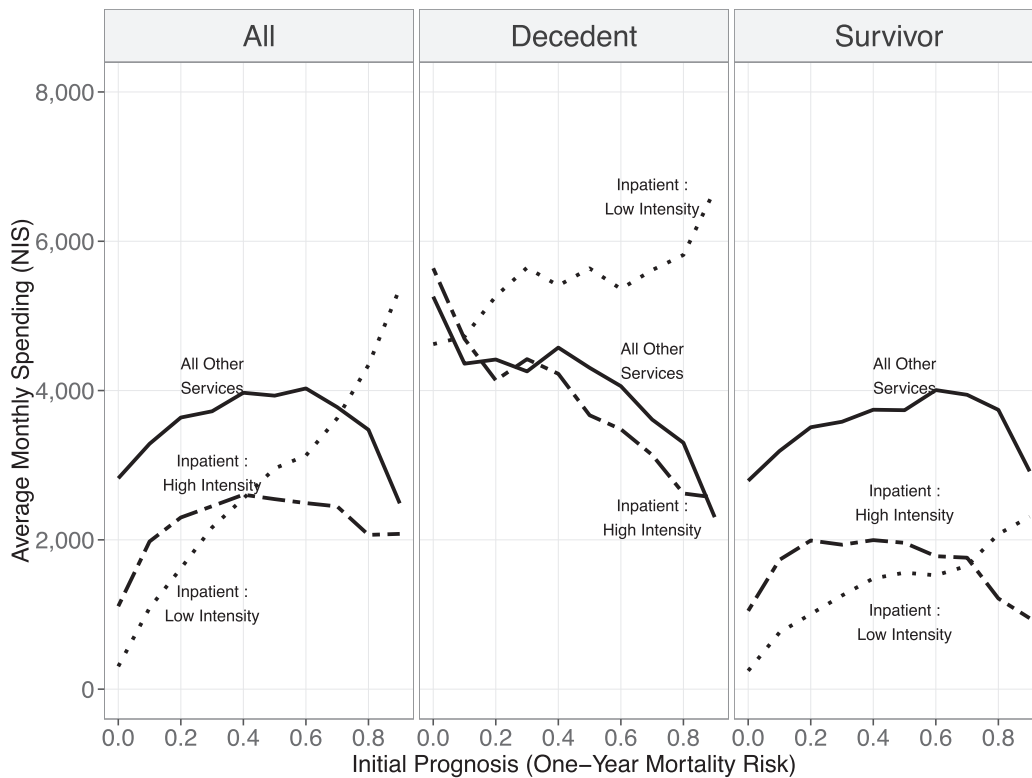
Types of services. Elevated spending for decedents is almost entirely driven by differences in inpatient spending (table 2).¹² Although inpatient spending only accounts for

40% of medical spending among survivors, higher spending on inpatient care accounts for 95% of the elevated spending on decedents. Spending on all other care, including outpatient care, radiation, and chemotherapy, is only 8% larger among decedents than among survivors with a similar initial prognosis.

Elevated inpatient spending in turn is disproportionately concentrated in low-intensity (versus high-intensity) admissions and in unplanned (versus planned) admissions. Despite accounting for only a quarter of inpatient spending among survivors, low-intensity admissions account for almost

¹²We focus our discussion on the comparison of decedents to survivors reweighted to have the same distribution of initial prognoses as decedents, so that we are comparing treatment of patients with the same initial prognosis.

FIGURE 4.—AVERAGE MONTHLY SPENDING, BY TYPE OF SERVICE AND INTENSITY



The figure shows average monthly spending (in the 12 months postdiagnosis) as a function of initial prognosis (one-year mortality risk), separately for low-intensity admissions, high-intensity admissions, and all other services. Panels show results separately for all patients (left), decedents (middle), and survivors (right). Decedent spending is adjusted for survival duration (see equation [2]). All spending measures are in current New Israeli Shekels (NIS). $N = 83,181$ patients.

two-thirds of the elevated spending on decedents. Likewise, unplanned admissions account for only about a quarter of inpatient spending among survivors, but about half of the elevated spending on decedents.¹³

Average monthly spending on low-intensity admissions is strongly increasing with poorer initial prognosis (figure 4). In other words, the poorer the patient's initial prognosis, the greater the spending on low-intensity admissions. By contrast, spending on high-intensity admissions and spending on other services are fairly flat or declining with initial prognosis.

Survivor-decedent differences in inpatient spending reflect differences in inpatient use (table 3). In the year following diagnosis, decedents are twice as likely to have a hospital admission each month: 41.8% compared to 21.5% of survivors for the same initial prognosis. Moreover, conditional on having an admission in a given month, decedents have 1.9 admissions per month, compared with 1.6 for survivors. Length of stay is also longer for decedents, on average 9.3 days per

admission, compared with 7.4 days for survivors. And as with hospital spending, decedent hospital utilization is also concentrated in low-intensity admissions. Every month in the year following initial diagnosis, decedents are 35% more likely to have a high-intensity admission but nearly *three times* more likely to have a low-intensity admission (31.8% of decedents compared to 12.5% for survivors). Moreover, conditional on having any admission, decedents have 0.36 additional low-intensity admissions and 0.12 *fewer* high-intensity admission.

Spending patterns over the course of treatment. Treating cancer is a dynamic process, typically consisting of a sequence of decisions, each depending on the results of earlier stages. We examine how changes in prognosis over the first year correlate with subsequent changes in spending and spending type. Because they are based on the sample of patients still in treatment, these results do not directly relate to the decedent-survivor difference in spending. Nonetheless, restricting attention to patients while they are in treatment provides an alternative perspective on the relationship of risk and spending. It complements the previous analyses and provides a window into the process of dying, while relying solely on information available in real time.

We find that a worsening of the prognosis is associated with an increase in subsequent spending. Figure 5 examines pairs

¹³An alternative way to classify admissions is based on whether they were billed as a procedure-based bundled episode or per-diem. Procedure-based billing is only appropriate for admissions with a major therapeutic procedure, such as surgery or inpatient chemotherapy. As a result, procedure-based billed admissions are more expensive than admissions billed per diem (appendix table A4). The concentration of spending on decedents in relatively low-cost admissions persists when we use this alternative measure to classify admission intensity (appendix table A5).

TABLE 3.—MONTHLY ADMISSION STATISTICS

	Survivor		Decedent	Difference
	Unweighted (1)	Reweighted by decedent risk (2)	(3)	Decedent-survivor (reweighted) (4)
A. Any admission				
All	11.4	21.5	41.8	20.3
Low intensity	4.2	12.5	31.8	19.3
High intensity	7.9	11.0	14.8	3.8
B. Admissions per month (if any during the month)				
All	1.5	1.6	1.9	0.2
Low intensity	0.6	1.0	1.4	0.4
High intensity	0.9	0.6	0.5	−0.1
C. Length of stay (days)				
All	6.0	7.4	9.3	1.8
Low intensity	6.4	7.1	8.8	1.7
High intensity	5.7	7.9	10.5	2.6

Table shows monthly admission statistics in the 12 months post-cancer diagnosis. Columns show results separately for survivors and decedents. Survivor statistics in column 2 are reweighted by decedent risk and month-from-diagnosis (see equation [3]). Decedent-survivor is the difference between Decedent and Survivor (reweighted) outcomes. In panel A, any admission shows the fraction of patients with an admission during each month over the period during which each patient was still alive during the first year after initial diagnosis. In panel B, admissions per month shows the average number of admissions for months during which the patient had at least one admission. In panel C, length of stay is the average duration of stay over all admissions. Within each panel, we partition admissions into low-intensity and high-intensity admissions, as described in the text. $N = 83,181$ patients.

of adjacent major clinical events and shows the relationship between the change in the current mortality prognosis and the change in subsequent average monthly spending between these events.¹⁴ On average, a 5 percentage point increase in mortality risk between events is associated with about a NIS 1,000 increase in subsequent average monthly spending, but the relationship is concave; greater increases in risk result in only slightly higher increases in spending (panel A). The association between the change in mortality prognosis and the change in spending also depends on the current *level* of predicted mortality risk: the worse the current prognosis, the weaker this association is (panel B).¹⁵ These results are consistent with treatment intensifying for complex cases that do not respond well to previous treatments. But such intensification is not without limits: care for cases with a grim prognosis does not intensify even when the prognosis further worsens.

We also examined what *types* of clinical events are associated with a worsening prognosis. The results show that low-intensity admissions—and only low-intensity admissions—are associated with a pronounced worsening of prognosis; that is, an increase in mortality risk (figure 6). On average, a low-intensity admission is associated with an increase of more than 10 percentage points in mortality risk. When such admissions start, it is not a good sign.

Figure 7 returns to the patient-level sample of all cancer decedents to explore these time patterns from another

perspective. In separate panels by type of service, it shows decedent spending as a function of two different timelines: months after diagnosis (panel A), and months before death (panel B). Each line shows the average monthly spending of a group of decedents who survived the same integer number of months. To the extent that services reflect treatment plans that are decided in advance, we would expect to see the timing of spending aligned on a prospective time scale (top panels), regardless of eventual survival duration. In contrast, treatment responses to unexpected deterioration may be better aligned with the retrospective time scale (bottom panels), regardless of survival duration. The results show that regardless of survival duration, low-intensity admissions spike in the last couple of months before death. In contrast, spending on high-intensity admissions and on other services (including outpatient services and drugs) spike two or three months after diagnosis and decreases in the last month or two.

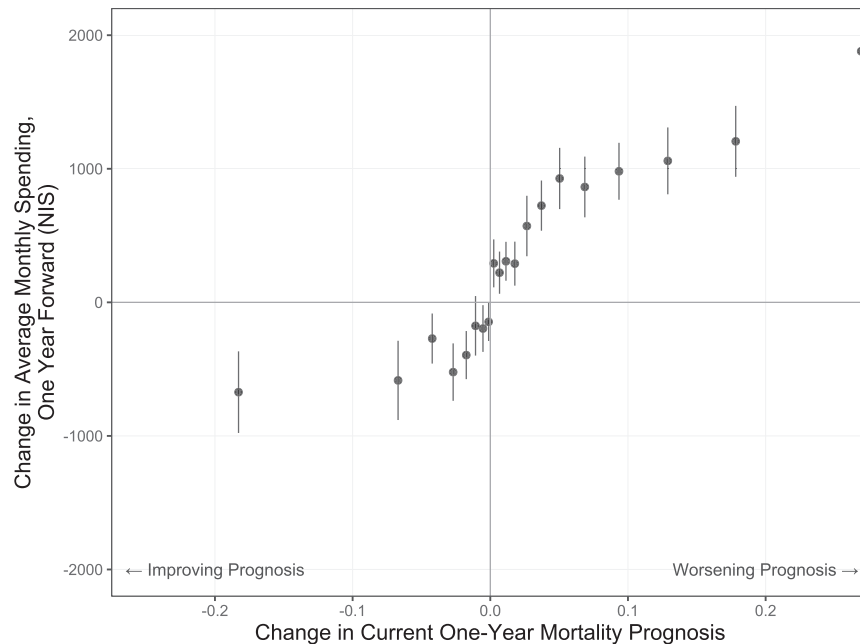
Table 4 summarizes these patterns quantitatively. Closer to death, decedent admissions involve fewer surgeries and more maintenance relative to both decedent admissions farther from death and survivor admissions. Overall, 27.6% of admissions for cancer patients involve surgery. However, only 9.4% of decedent admissions in the last month before death involve surgery, compared with 11.2% of decedent admissions that occur four to 12 months before death, and with 33.4% of survivor admissions. Admissions closer to death also involve fewer chemotherapy procedures, more diagnostics, and more maintenance. Radiation does not have a clear trend (possibly because there are both therapeutic and palliative radiation therapies).

Overall, the results paint a reasonably clear picture in which the timing of high-intensity admissions and other services is primarily tied to the timing of cancer diagnosis, while the timing of low-intensity admissions is closely linked to the (retrospective) timing of death. This is consistent with initial treatment plans that fight cancer via scheduled surgeries, outpatient radiation, and chemotherapy but change to a different


¹⁴Subsequent spending is measured over a one-year period following each event and is adjusted for survival duration. We exclude from the measure of subsequent spending all spending associated with the current event; including such spending makes the relationship between deteriorating prognosis and increasing spending even stronger.

¹⁵We find similar results when instead of evaluating the heterogeneity in the association across different ranges of current mortality prognosis, as in panel B, we instead evaluate the heterogeneity across different quintiles of current prognosis, where quintiles are calculated within cancer type (appendix figure A6). This suggests that both absolute and relative risk levels mediate the association between the change in mortality prognosis and the change in spending.

FIGURE 5.—RELATIONSHIP BETWEEN CHANGE IN PROGNOSIS AND CHANGE IN SUBSEQUENT SPENDING



(A) Overall



(B) by Current Prognosis

Figure shows, for the sample of 207,607 clinical histories of cancer patients in our sample with one more clinical event after initial diagnosis, the relationship between the change in current prognosis and the change in forward spending, overall (panel A) and by level of current mortality prognosis (panel B). Each observation in the underlying data is a pair of consecutive clinical events. The x axis shows the change in mortality prognosis between the start of the most recent and the start of the current clinical event. The y axis shows the change in one-year forward spending between the two events. Linear fit is shown on panel B.

type of medical treatment if treatment has failed. This change involves an increased frequency of unplanned admissions that may aim to monitor and maintain patients without necessarily trying to treat them.

Of course, analyses that look back from the time of death are conducted from an *ex-post* perspective. It would be a mistake to conclude that because low-intensity admissions tend to spike close to the time of death—regardless of initial prognosis or survival time—reducing such events would reduce

spending without any harm to patients. For this to be the case, we would need to be able to predict, at the time of the admission, that these admissions are very likely at the end of life. Figure 8 shows that we cannot. It looks at the fraction of low-intensity and high-intensity admissions that result in death within 60 days, as a function of current prognosis at the time of admission. Admissions that result in near-term death rise sharply as the current prognosis worsens. However, many admissions do not result in near-term death, even among

FIGURE 6.—CHANGE IN CURRENT PROGNOSIS BETWEEN PREVIOUS AND CURRENT CLINICAL EVENTS, BY CURRENT EVENT TYPE

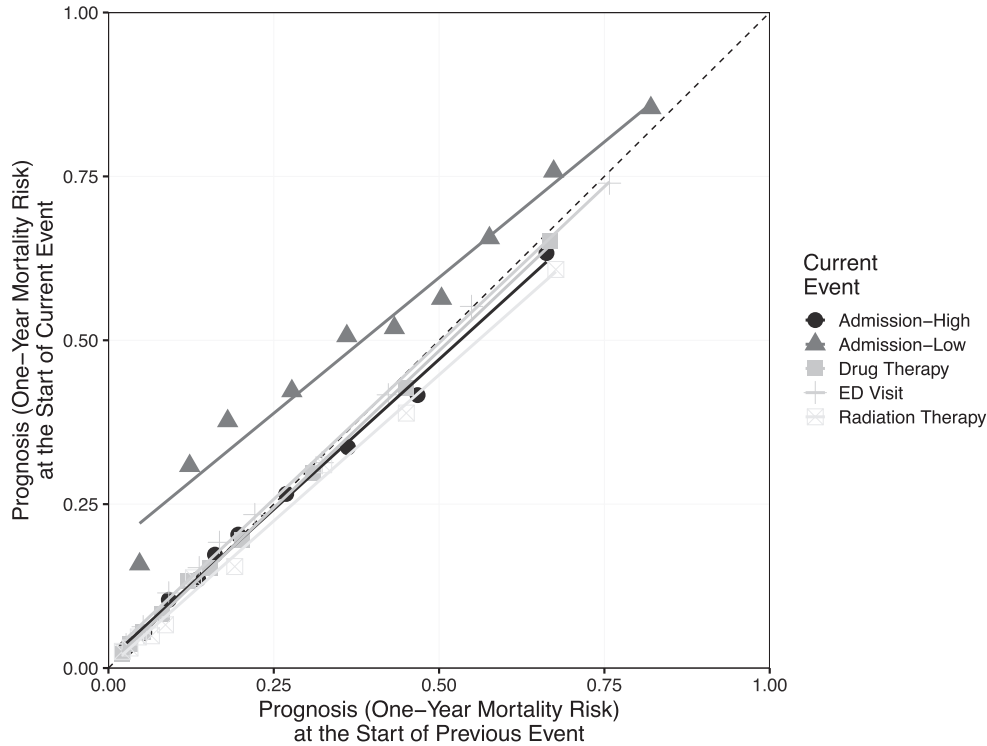


Figure shows, for the sample of 207,607 clinical histories of cancer patients in our sample with one more clinical event after initial diagnosis, the relationship between mortality prognosis at the start of the current clinical event over the mortality prognosis at the start of the previous event. Each observation in the underlying data represents a pair of consecutive clinical events, for patients with the same cancer type who had the same prior sequence of events. The x and y axes show the predicted mortality prognosis at the start of the previous and the current clinical event, respectively. Shape and color denote current event type. The data are binned by deciles of the previous mortality prognosis, separately for each (current) event type. Linear fit is shown for each risk group separately. The dashed line is the identity (45°) line. The underlying sample sizes, by current event type, are as follows: 69,745 high-intensity admissions, 43,897 low-intensity admissions, 48,089 drug therapy events, 12,631 radiation therapy events, and 33,245 ED visits.

TABLE 4.—INPATIENT PROCEDURES BY ADMISSION TIME BEFORE DEATH

	Procedure type, admission with any (%)						N of admissions (7)
	Maintenance (1)	Diagnostics (2)	Surgery (3)	Radiation (4)	Chemotherapy (5)	Other (6)	
Last month	11.5	98.7	10.0	4.0	5.1	0.7	5,219
1–3 months	11.9	96.0	12.0	6.5	9.5	0.9	3,864
4–12 months	11.0	94.0	16.3	6.6	16.7	1.5	4,989
Survivors	9.2	90.2	34.0	2.9	7.7	1.1	36,596
All	9.8	91.9	28.1	3.7	8.4	1.1	50,668

The fraction of sampled admissions that included procedures of different types. Sampled admissions include Clalit-owned-hospital admissions that started and ended during the year after diagnosis. Appendix table A6 shows data separately for unplanned and planned admissions and for high- and low-intensity admissions.

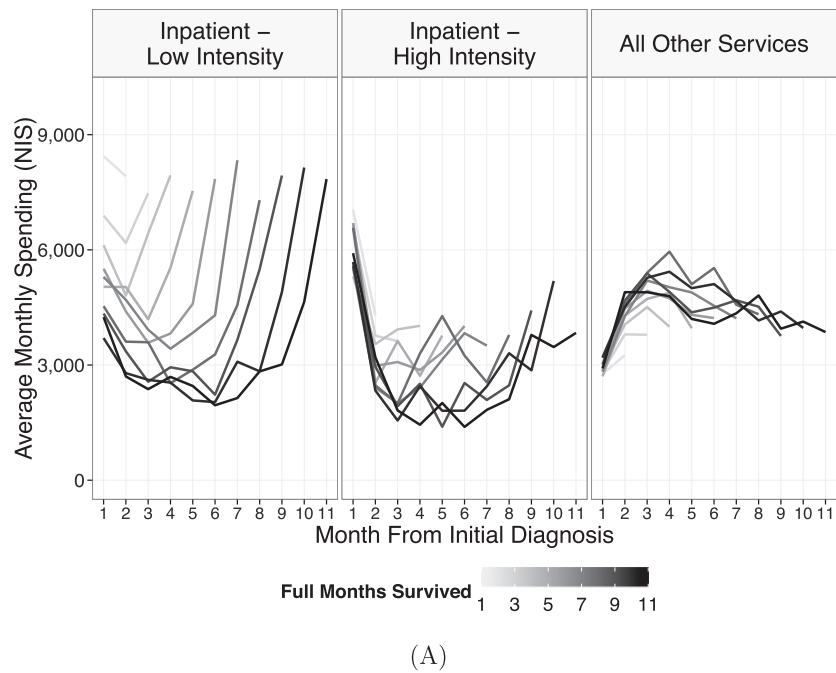
individuals with poor current prognoses. For example, among patients who enter a low-intensity admission with a current prognosis of 80% mortality within a year, less than half die within the next two months.

Elevated spending on decedents, by age. Cancer is a disease that (unfortunately) affects a wide range of ages. We can therefore examine how the elevation of spending on decedents varies by age. Among patients with the same initial prognosis, average monthly spending declines with age; this decline is particularly pronounced for decedents compared to survivors (figure 9). Table 5 summarizes these results quantitatively and shows that the difference in average monthly spending for decedents, relative to survivors with the same initial prognosis, decreases monotonically with age. The el-

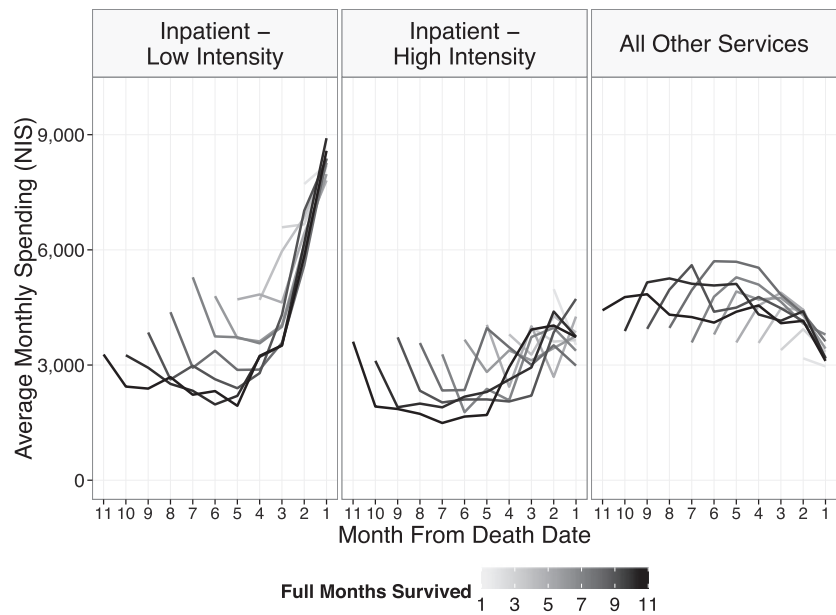
evaluation of spending on decedents is about NIS 9,500 for the youngest age quintile (53 years old or younger) but declines to about NIS 4,000 for the highest age quintile (78 years old and older).¹⁶ This pattern persists if we look within cancer topography (panel B). For example, among breast cancer patients, the difference in spending between decedents and survivors is NIS 5,500 for the youngest age quintile and NIS 3,600 for the oldest; for stomach cancer patients, the difference for the youngest and oldest age quintiles is NIS 5,500 and 2,500, respectively.

¹⁶These differences do not reflect age differences in survival duration: additional analyses (not shown) confirm that conditional on the initial prognosis, decedents younger and older than the median for their cancer type have the same expected survival duration.

FIGURE 7.—DECEDENT SPENDING BY TIME BEFORE DEATH AND AFTER DIAGNOSIS



(A)



(B)

Panels show average monthly spending data for the sample of cancer decedents ($N = 16,289$). In both panels, each line represents average spending for a group of decedents who survived the same integer number of months, excluding partial months' spending, with darker lines representing longer survival. However, in panel A, the horizontal axis counts the number of months from the index date, whereas in panel B, the horizontal axis counts the number of months before death. In both cases, we show results separately for low-intensity inpatient admissions, high-intensity inpatient admissions, and all other services. All spending measures are in current New Israeli Shekels (NIS).

To shed more light on the source of the age gradient, we examine how spending varies by age as a function of *current* mortality prognosis. The results are revealing. Conditional on current prognosis and the type of current episode, spending on that episode is similar for the old and young (figure 10, panel A). By contrast, average adjusted monthly spending in the year after the current episode is substantially more pro-

nounced on the young (panel B). This suggests that the age differences in overall spending is driven not by differences in the cost of specific episodes, but by younger patients receiving more therapies, and (as seen in appendix figure A7) more intensive therapies over the course of their treatment.

These age patterns are striking, although their interpretation is not obvious. The results are not consistent with an

FIGURE 8.—FRACTION OF ADMISSIONS ENDING IN DEATH WITHIN 60 DAYS, BY CURRENT PREDICTED MORTALITY

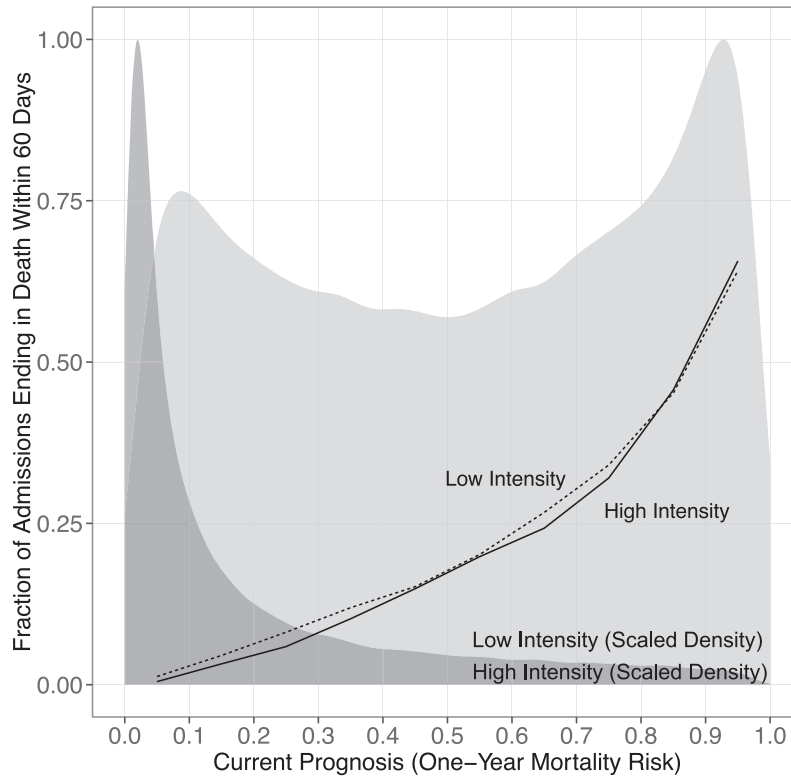


Figure shows the fraction of admissions ending in death within 60 days of admission, as a function of current prognosis (one-year mortality risk), as predicted at the first day of the admission. Results are shown separately for high-intensity and low-intensity admissions. Shaded areas are scaled densities of the current prognosis for high- and low-intensity admissions. The sample includes all admissions of cancer patients in the first year after diagnosis.

FIGURE 9.—SPENDING AND MORTALITY OF DECEDENTS AND SURVIVORS, BY AGE QUINTILES

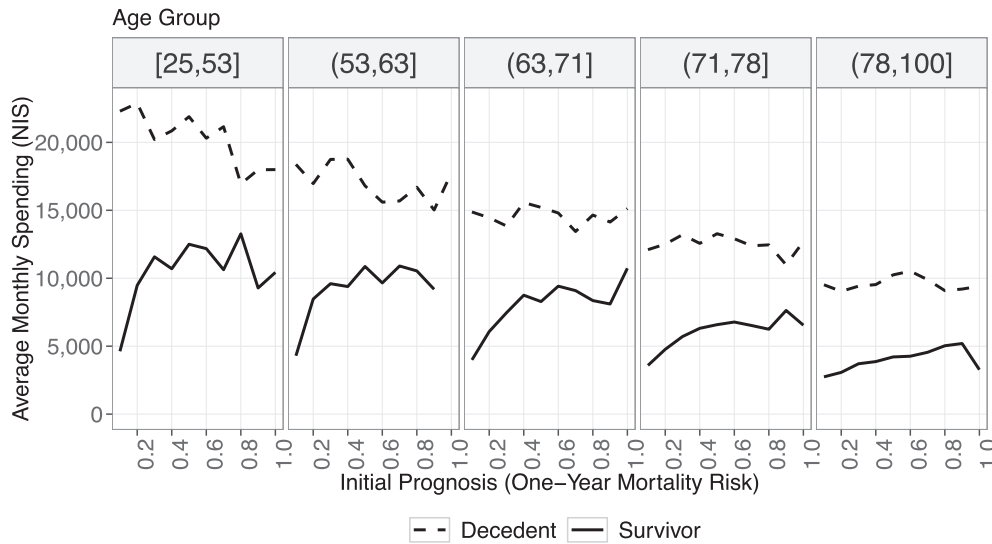


Figure shows, separately by age quintiles, average monthly spending on all services by initial prognosis (one-year mortality risk). Decedent spending is adjusted for survival duration (see equation [2]). The top quintile is top-coded at 100 years of age. All spending measures are in current New Israeli Shekels (NIS). $N = 83,181$ patients.

explanation based on differences by age in the expensiveness of a given treatment. One natural possibility is that appropriateness of treatments varies with age. By conditioning on cancer type in some of the analyses we tried, to the extent possible, to hold fixed the available “technology” or treat-

ment options, although of course there may be remaining differences by age.

These patterns suggest that treatment decisions may be made not only with respect to their likely short term (i.e., within a year) impact, but also factoring in the longer life

TABLE 5.—AVERAGE MONTHLY SPENDING, BY AGE QUINTILE

	Age quintile (1)	Survivor		Decedent (4)	Difference
		Unweighted (2)	Rewighted by decedent risk (3)		Decedent-survivor (reweighted) (5)
A. All cancer types					
	[25,53]	5,308	10,943	20,484	9,541
	(53,63]	5,141	10,010	16,760	6,750
	(63,71]	4,754	8,964	14,297	5,333
	(71,78]	4,256	7,044	12,252	5,208
	(78,100]	3,311	5,354	9,384	4,030
B. By cancer type					
Breast	[25,53]	6,793	7,073	12,561	5,487
	(78,100]	2,441	2,834	6,486	3,651
Prostate	(53,63] [*]	2,713	2,954	10,629	7,675
	(78,100]	2,603	2,662	6,891	4,229
Colon	[25,53]	6,152	8,880	18,269	9,389
	(78,100]	3,679	5,675	9,393	3,718
Bronchus and lung	[25,53]	7,748	11,244	14,686	3,442
	(78,100]	4,680	5,874	9,228	3,355
Skin	(53,63] [*]	1,475	1,604	16,525	14,921
	(78,100]	1,690	2,461	6,988	4,527
Bladder	(53,63] [*]	2,376	3,622	14,419	10,798
	(78,100]	2,565	3,411	10,181	6,770
Hematopoietic system	[25,53]	15,038	22,025	47,886	25,861
	(78,100]	3,762	5,113	9,843	4,731
Lymph nodes	[25,53]	9,363	13,913	31,017	17,104
	(78,100]	6,938	10,216	12,870	2,654
Stomach	[25,53]	6,496	11,451	17,033	5,582
	(78,100]	4,533	6,608	9,071	2,463

Table shows average monthly spending in the 12 months post-cancer diagnosis for different age groups, by quintiles of patient age at the time of cancer diagnosis. Column 1 shows the age range, with square brackets and parentheses denoting included and excluded endpoints, respectively. Columns 2–4 show results separately for decedents and survivors. Decedent spending is adjusted for survival duration (see equation [2]). Survivor spending in column 2 is reweighted by decedent risk and month-from-diagnosis (see equation [3]). Decedent-Survivor (column 5) is the difference between Decedent and Survivor (reweighted) spending. All spending measures are in current New Israeli Shekels (NIS). Panel A shows results for all cancer types, by patient age quintile. Panel B shows results for youngest and oldest age quintiles, for the most common cancer types in our sample. For cases marked by *, the youngest age group [25,53] did not have sufficiently many decedents in all bins for reweighting, so the second-youngest age group (53,63] is shown instead. $N = 83,181$ patients.

expectancy of the young, conditional on successful treatment. In other words, the “return on investment” is higher for younger patients.¹⁷ In addition, the willingness to let go or the psychological cost of conceding defeat may well be higher for younger patients. Whatever the underlying mechanism, these results are suggestive evidence against a narrative of patients (and their doctors and families) making treatment decisions with an exclusive focus on near-term survival.

V. Conclusion

We looked inside the “black box” of elevated spending at the end of life, using an extremely rich dataset on a large population and focusing on newly diagnosed cancer patients, who represent a relatively homogeneous set of medical conditions. We have three main findings concerning the sources of elevated spending on decedents relative to survivors with the same initial prognosis. The elevated spending for decedents is almost entirely driven by inpatient spending (particularly low-intensity admissions with few procedures). A worsening prognosis is strongly associated with an increase


¹⁷Of course, the social benefits of end-of-life spending may be greater than their individual benefits if the use of new treatment generates positive externalities to other patients by facilitating learning about the efficacy of such treatments.

in low-intensity admissions—in other words, with (expensive) inpatient maintenance care. Finally, the concentration of spending among decedents is particularly pronounced for younger cancer patients relative to older cancer patients.


In addition to a “forensic accounting” exercise for end of life spending, our descriptive analyses raise several potential implications. In contrast to the hypothesis that treatment decisions are made from a very present-focused goal of prolonging life over the very near term, the age-related findings are consistent with decisions being made at least partly with a longer-term horizon in mind. In particular, the higher elevation of spending on decedents (relative to survivors with similar initial prognoses) for younger patients compared to older patients is consistent with greater demand (among patients, their families, and their physicians) for treating those with a higher life expectancy (conditional on surviving cancer). These patterns may also point to a role for preferences—perhaps a greater reluctance to “let go” among the young—to influence end-of-life spending patterns. Such reluctance may also explain prior evidence that healthcare spending on pets spikes at the end of life as well (Einav et al., 2017).

Another important implication of our findings arises from what we did not find: the dog that did not bark. Specifically, our results repeatedly stop short of identifying any clear

FIGURE 10.—SPENDING ON CURRENT AND FUTURE EVENTS, BY CURRENT PROGNOSIS AND AGE



(A) Spending on Current Event



(B) Spending One Year Forward (or Until Death)

Figure shows, using the sample for which we predict current mortality based on major clinical events ($N = 292,284$ patient-events), the average current and future spending over current mortality prognosis, by age. Old (Young) denotes patients whose age at the time of cancer diagnosis is older (younger) than the median for their cancer type. Current prognosis (one-year mortality risk) is predicted at the start of the current event. Facets show data separately by current event type; the rightmost facet shows data for all event types combined. Panel A shows spending on the current event. Panel B shows spending one-year forward, excluding the current event. Both current and forward spending is measured as a monthly average (adjusted for survival duration based on equation [2]), so results presented in the two panels are comparable. See section III for detailed definitions.

categories of spending that could be reduced without concern about potential patient harm. Even among low-intensity admissions for patients with very poor current prognoses, a substantial share of admissions do not result in near-term death. This underscores the perennial challenge of identifying “obvious” ways to reduce large amounts of healthcare spending. A more fruitful (although also more laborious) path to identifying waste in healthcare systems may lie in credibly documenting the many specific, smaller sources of spending that could be eliminated with little or no harm to patients, as recent research has started to do (Abaluck et al., 2016;

Einav, Finkelstein, & Mahoney, 2019; Cooper, Scott Morton, & Shekita, 2019).

REFERENCES

- Abaluck, Jason, Leila Agha, Chris Kabrbel, Ali Raja, and Arjun Venkatesh, “The Determinants of Productivity in Medical Testing: Intensity and Allocation of Care,” *American Economic Review* 106:12 (2016), 3730–3764. 10.1257/aer.20140260
- Barnato, Amber E., Mark B. McClellan, Christopher R. Kagay, and Alan M. Garber, “Trends in Inpatient Treatment Intensity Among Medicare Beneficiaries at the End of Life,” *Health Services Research* 39:2 (2004), 363–376. 10.1111/j.1475-6773.2004.00232.x, PubMed: 15032959

- Bekelman, Justin E., Scott D. Halpern, Carl Rudolf Blankart, Julie P. Bynum, Joachim Cohen, Robert Fowler, Stein Kaasa, Lukas Kwietniewski, Hans Olav Melberg, Bregje Onwuteaka-Philipsen et al., "Comparison of Site of Death, Health Care Utilization, and Hospital Expenditures for Patients Dying With Cancer in 7 Developed Countries," *JAMA* 315:3 (2016), 272–283. 10.1001/jama.2015.18603, PubMed: 26784775
- Chen, Tianqi, and Carlos Guestrin, "Xgboost: A Scalable Tree Boosting System," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (New York: ACM, 2016, pp. 785–794).
- Cooper, Zack, Fiona Scott Morton, and Nathan Shekita, "Surprise! Out-of-Network Billing for Emergency Care in the United States," NBER working paper 23623, 2019.
- DeSalvo, Karen B., Vincent S. Fan, Mary B. McDonnell, and Stephan D. Fihn, "Predicting Mortality and Healthcare Utilization With a Single Question," *Health Services Research* 40:4 (2005), 1234–1246. 10.1111/j.1475-6773.2005.00404.x, PubMed: 16033502
- Einav, Liran, Amy Finkelstein, and Atul Gupta, "Is American Pet Health Care (Also) Uniquely Inefficient?," *American Economic Review* 107:5 (2017), 491–95. 10.1257/aer.p20171087
- Einav, Liran, Amy Finkelstein, and Neale Mahoney, "Long-Term Care Hospitals: A Case Study in Waste," NBER working paper 24946, 2019.
- Einav, Liran, Amy Finkelstein, Sendhil Mullainathan, and Ziad Obermeyer, "Predictive Modeling of US Health Care Spending in Late Life," *Science* 360:6396 (2018), 1462–1465. 10.1126/science.aar5045, PubMed: 29954980
- Emanuel, Ezekiel J., and Linda L. Emanuel, "The Economics of Dying—The Illusion of Cost Savings at the End of Life," *New England Journal of Medicine* 330:8 (1994), 540–544. 10.1056/NEJM199402243300806, PubMed: 8302321
- Emanuel, Ezekiel J., Arlene Ash, Wei Yu, Gail Gazelle, Norman G. Levinsky, Olga Saynina, Mark McClellan, and Mark Moskowitz, "Managed Care, Hospice Use, Site of Death, and Medical Expenditures in the Last Year of Life," *Archives of Internal Medicine* 162:15 (2002), 1722–1728. 10.1001/archinte.162.15.1722, PubMed: 12153375
- Fischhoff, Baruch, "Hindsight is Not Equal to Foresight: The Effect of Outcome Knowledge on Judgment Under Uncertainty," *Journal of Experimental Psychology: Human Perception and Performance* 1:3 (1975), 288. 10.1037/0096-1523.1.3.288
- French, Eric B., Jeremy McCauley, Maria Aragon, Pieter Bakx, Martin Chalkley, Stacey H. Chen, Bent J. Christensen, Hongwei Chuang, Aurelie Côté-Sergent, Mariacristina De Nardi, et al., "End-of-life Medical Spending in Last Twelve Months of Life is Lower Than Previously Reported," *Health Affairs* 36:7 (2017), 1211–1217. 10.1377/hlthaff.2017.0174, PubMed: 28679807
- Genevès, Pierre, Thomas Calmant, Nabil Layaïda, Marion Lepelley, Svetlana Artemova, and Jean-Luc Bosson, "Predicting At-Risk Patient Profiles from Big Prescription Data," *fhal-01517087v4f [Reprint]* (2017).
- Heron, Melonie, "Deaths: Leading Causes for 2010," *National Vital Statistics Reports, National Center for Health Statistics, Center for Disease Control and Prevention* 62:6 (2013), 1–96.
- Makar, Maggie, Marzyeh Ghassemi, David M. Cutler, and Ziad Obermeyer, "Short-term Mortality Prediction for Elderly Patients Using Medicare Claims Data," *International Journal of Machine Learning and Computing* 5:3 (2015), 192. 10.7763/IJMLC.2015.V5.506
- Medicare Payment Advisory Commission, "Report to the Congress: Improving Care at the End of Life," Technical Report (1999).
- Morden, Nancy E., Chiang-Hua Chang, Joseph O. Jacobson, Ethan M. Berke, Julie P. W. Bynum, Kimberly M. Murray, and David C. Goodman, "End-of-Life Care for Medicare Beneficiaries With Cancer is Highly Intensive Overall and Varies Widely," *Health Affairs* 31:4 (2012), 786–796. 10.1377/hlthaff.2011.0650, PubMed: 22492896
- Nicholas, Lauren Hersch, Kenneth M. Langa, Theodore J. Iwashyna, and David R. Weir, "Regional Variation in the Association Between Advance Directives and End-of-Life Medicare Expenditures," *JAMA* 306:13 (2011), 1447–1453. 10.1001/jama.2011.1410, PubMed: 21972306
- Riley, Gerald F., and James D. Lubitz, "Long-Term Trends in Medicare Payments in the Last Year of Life," *Health Services Research* 45:2 (2010), 565–576. 10.1111/j.1475-6773.2010.01082.x, PubMed: 20148984
- Teno, Joan M., Pedro L. Gozalo, Julie P. W. Bynum, Natalie E. Leland, Susan C. Miller, Nancy E. Morden, Thomas Scupp, David C. Goodman, and Vincent Mor, "Change in End-of-Life Care for Medicare Beneficiaries: Site of Death, Place of Care, and Health Care Transitions in 2000, 2005, and 2009," *JAMA* 309:5 (2013), 470–477. 10.1001/jama.2012.207624, PubMed: 23385273
- Zeltzer, Dan, Ran D. Balicer, Tzvi Shir, Natalie Flaks-Manov, Liran Einav, and Efrat Shadmi, "Prediction Accuracy with Electronic Medical Records Versus Administrative Claims," *Medical Care* 57:7 (2019), 551–559. 10.1097/MLR.0000000000001135, PubMed: 31135691