American Economic Review 2023, 113(11): 3003–3043
https://doi.org/10.1257/aer.20211638

Is There a VA Advantage? Evidence from Dually Eligible Veterans†

By David C. Chan, David Card, and Lowell Taylor*

We study public versus private provision of health care for veterans aged 65 and older who may receive care provided by the US Department of Veterans Affairs (VA) and in private hospitals financed by Medicare. Utilizing the ambulance design of Doyle et al. (2015), we find that the VA reduces 28-day mortality by 46 percent (4.5 percentage points) and that these survival gains are persistent. The VA also reduces 28-day spending by 21 percent and delivers strikingly different reported services relative to private hospitals. We find suggestive evidence of complementarities between continuity of care, health IT, and integrated care. (*JEL* H51, I11, I12, I18, J14)

A key question in the design of health care systems worldwide is whether the government or the private sector should provide care. In the United States, the choice between public and private provision has become a top policy issue for the Department of Veterans Affairs (VA). Seeking to improve veteran access to health care, policymakers have debated whether the VA should expand the capacity of its system—the Veterans Health Administration—or shift health care delivery to private providers.

An extensive descriptive literature (e.g., Reid 2010; Blank, Burau, and Kuhlmann 2017) has compared health care outcomes in public versus private systems. More generally, economists have long debated the appropriate size and role of the public sector in the economy, highlighting theoretical arguments about competitive pressure, ownership structure, and differences in the objectives and constraints in the public versus private sectors (Alchian 1965; Stigler 1965; Hart, Shleifer, and Vishny 1997). Nevertheless, rigorous empirical evaluations of the

*Chan: Stanford University, Department of Veterans Affairs, and NBER (email: david.c.chan@stanford.edu); Card: UC Berkeley and NBER (email: card@berkeley.edu); Taylor: Carnegie Mellon University and NBER (email: lt20@andrew.cmu.edu). Chinhui Juhn was the coeditor for this article. We thank Steve Asch, Zarek Brot-Goldberg, Marika Cabral, Joe Doyle, Mark Duggan, Amy Finkelstein, Josh Gottlieb, Kate Ho, Peter Hull, Bob Kaestner, Chad Kessler, Tim Layton, Carol Propper, Michael McWilliams, Maria Polyakova, Jim Rebitzer, Adam Sacarny, David Silver, Anita Vashi, Todd Wagner, and Chris Walters for helpful comments and suggestions. We also thank Sophie Andrews, Noah Boden-Gologorsky, Sydney Costantini, Kaveh Danesh, Akriti Dureja, Johnny Huynh, Kevin Kloiber, Uyseok Lee, Chris Lim, Matthew Merrigan, Kevin Todd, Justine Weng, and Saam Zahedian for excellent research assistance. Finally, we thank Melissa Taylor for valuable insight concerning our research design. Chan gratefully acknowledges generous support from VA HSR&D I01HX002631, the Laura and John Arnold Foundation, and the Stanford Institute for Economic Policy Research. This study was approved by the Institutional Review Board at Stanford University (eProtocol #34898) and by the Scientific Review Subcommittee at the VA Palo Alto Health Care System (RDIS #CHD005).The contents of this paper do not represent the views of the VA or the United States government.

†Go to https://doi.org/10.1257/aer.20211638 to visit the article page for additional materials and author disclosure statements.

performance of public versus private health care providers have been rare. Public and private providers usually serve different patient populations, either by statute or by patient selection.

In this paper, we focus on "dually eligible" veterans aged 65 and older who can receive health care at both VA facilities and private hospitals that accept Medicare. We use the ambulance design proposed by Doyle et al. (2015) to study the causal effect of receiving emergency care at the VA versus a non-VA facility. Our approach compares veterans sharing key characteristics—zip code of residence, prior VA and non-VA health care utilization, and location of pickup (e.g., home residence, nursing home)—who receive the same dispatched level of ambulance service (i.e., advanced versus basic life support) from different ambulance companies. Our main analytic sample includes 401,319 911-dispatched ambulance rides from 2001 to 2014 for veterans with prior attachment to the VA and in a zip code served by at least two ambulance companies. As in Doyle et al. (2015), we show that the leave-out share of dually eligible veterans transported to the VA by the assigned ambulance company is a strong predictor of hospital assignment. Under the plausible assumption that ambulances are quasi-randomly assigned within zip codes and in cells of key characteristics, this design allows us to study the effect of VA versus non-VA emergency care on health outcomes.

We find that in the high-mortality population of elderly veterans with emergencies, there is a VA advantage—a 46 percent reduction in 28-day mortality relative to the baseline (4.5 pp, with a 95 percent confidence interval of 1.1 to 8.0 pp). We show that our instrumental variables (IV) estimates of the VA effect are robust to controlling for a long list of characteristics of both the index patient and other patients transported by the same ambulance company. The latter set of ambulance co-rider controls can account for unobserved selection patterns across ambulance companies (Altonji and Mansfield 2018). The IV estimates are larger in magnitude than the corresponding OLS estimates, which center around 2.4 pp, with tight confidence intervals. A possible explanation for this difference is that VA "always-takers" (patients taken to the VA even by ambulance companies with the lowest VA rates) have worse health than VA "never-takers" (those taken to private hospitals even by ambulance companies with the highest VA rates). This selection pattern has been suggested by the medical literature; we examine it in greater detail below (Agha et al. 2000).

An important question for interpreting the survival benefits of VA care is whether these effects fade over longer horizons—as would happen if VA emergency care only temporarily displaces the mortality of fragile patients under "harvesting" (Schwartz 2000). To address this, we use an insight from Abadie (2002) to estimate the weekly potential death rates in the year after the initial ambulance ride among compliers of the quasi-experiment, i.e., patients whose destination hospital is determined by the ambulance company. With this tool, we disentangle the short-term versus long-term effects of the VA. Despite a high long-term mortality rate (close to one in three veterans will be dead within one year of the ambulance ride), we find that the mortality impact of presenting at the VA is concentrated in the first week, suggesting VA survival gains from care addressing temporary emergency conditions. We find no evidence of harvesting; the survival gains appear to be long lasting. Relying on intuition from Kitagawa (2015), we also use this potential outcomes framework to develop a sharper test of IV validity than the tests typically found in the applied literature.[1] Finally, we use this framework to document and account for widening differences between OLS and IV estimates of the VA advantage over longer time horizons. We interpret these widening differences as evidence of *higher* long-run mortality hazards for VA always-takers than for VA never-takers, strongly suggesting that veterans who select into the VA are indeed sicker.

The key potential threat to our research design is the possibility that, within zip codes, veterans assigned to ambulance companies more likely to go to the VA are healthier than those assigned to companies less likely to go to the VA. We present three additional pieces of evidence to address this concern. First, we show balance in characteristics of patients assigned to companies with different propensities of taking patients to the VA. Second, we conduct an extensive analysis along the lines suggested by Altonji, Elder, and Taber (2005), evaluating the stability of our estimates as we add controls to the models, including controls that measure the characteristics of *other* patients transported by the company. Third, in heterogeneity analyses, we show that the VA advantage is highly stable across VA and non-VA hospital characteristics that may be related to patient selection.

In the final section of the paper, we evaluate the mechanisms behind the VA advantage. First, we assess heterogeneity in the VA advantage according to patient and hospital characteristics. This heterogeneity could imply VA specialization in types of care needed most by veterans; it could also imply an advantage for medically vulnerable patients arising from continuity of care at the VA. Second, along the lines of Doyle et al. (2015), the VA could achieve better outcomes by spending more or by offering different services. Third, the VA advantage may reflect better access to patient information and coordination of care, particularly in high-uncertainty and high-stakes environments such as emergency care. This last mechanism is consistent with literature that highlights integration of care and health information technology (IT) as distinguishing features of the VA (McCarthy and Blumenthal 2006; Jha et al. 2009).

In the first class of mechanisms, we find evidence for moderate selection on gains. Compliers in our quasi-experiment—veterans with greater prior VA utilization who are medically needier and socioeconomically disadvantaged—tend to have higher treatment effects. However, we do not find that the VA worsens health outcomes in any set of patients or any location. While the VA hospitals differ from non-VA hospitals in their characteristics (e.g., they are more likely to be teaching hospitals), we also find a consistent VA advantage regardless of VA or non-VA hospital characteristics. This evidence suggests a widespread VA advantage, though disadvantaged veterans with complex medical needs appear to benefit more than other veterans from the VA.

[1] Specifically, we use the fact that, under IV validity, all indicators for potential outcomes must occur with positive probability among compliers (Balke and Pearl 1997; Kitagawa 2015). In the survival setting, this implies that the incremental mortality risk must be positive for compliers every week after the ambulance ride, both among those assigned to VA and non-VA hospitals. This prediction may fail if monotonicity violations arise, for example, because ambulance companies with higher VA propensities are *less* likely to send veterans with certain potential mortality outcomes to the VA. Chan, Gentzkow, and Yu (2022) show that this approach may detect violations in IV validity that remain hidden under standard "judges design" tests of monotonicity (e.g., Arnold, Dobbie, and Yang 2018).

We evaluate the second explanation by examining VA and Medicare spending, using information on actual spending by taxpayers and veterans. Spending following VA care is *lower* by $2,598, or about 21 percent, at 28 days. This suggests that the VA is more productive, achieving better outcomes at lower cost. We generate these results on actual spending from the perspective of taxpayers and patients. Alternatively, if we measure resource utilization by applying the same Medicare prices to reported procedures in both VA and non-VA settings (Finkelstein, Gentzkow, and Williams 2016), the reduction in spending following transport to a VA hospital doubles. We find striking differences in reported utilization of specific services between the two settings. Some portion of these differences likely reflects "upcoding"—in which cases are more likely to be coded as complex (Dafny 2005; Geruso and Layton 2020)—in the Medicare setting relative to the VA. Yet many services that are reimbursed little by Medicare (e.g., telephone calls) are much more likely to be documented in the provision of care for veterans arriving at the VA. Thus, it is also plausible that actual care substantively differs between the two settings, where differences in payment systems may imply differences in objectives.

The third explanation centers on the idea that coordination and continuity of care in an integrated delivery system may improve health outcomes—an explanation consistent with the larger impacts of the VA on medically needy patients and those with greater prior attachment. Unfortunately, as is the case with the previous literature, we cannot show direct quasi-experimental evidence of this joint mechanism among veterans who use the VA: The VA's transition to integrated care predates the period for which data are available for analysis, and veterans with no prior attachment to the VA are seldom transported to the VA. Instead, we study this indirectly among veterans who only use *non-VA* hospitals in the context of two policy reforms intended to stimulate health IT and integrated care in the private sector.

Specifically, through a parallel ambulance quasi-experiment, we ask whether non-VA users benefit from being assigned to their most visited (non-VA) hospital in the year prior to the ambulance ride (i.e., their "modal" hospital) and whether the benefit of assignment to the modal hospital is affected by two reforms: (i) the Health Information Technology for Economic and Clinical Health (HITECH) Act of 2009 and (ii) legislation to spur participation in "Accountable Care Organizations" (ACOs) beginning in 2011 (Blumenthal 2010; Greaney 2011). We find that the modal hospital survival benefit increases from a negligible effect before the policy reforms to about 1.9 pp—approximately one-half of the VA survival benefit—after 2010. We also find tentative evidence linking the increase in the modal hospital survival benefit to hospital-specific dates of health IT adoption and, to a lesser degree, ACO participation.

Our findings contribute to three sets of related literature. First, the public versus private provision of health care is a central question for the field of comparative health policy (Blank, Burau, and Kuhlmann 2017). The literature in this field has been mainly descriptive, comparing health care systems across the world.[2] Such comparisons are intrinsically difficult for the obvious reason that countries

[2] As an example of the amount of material devoted to such comparative studies, the European Observatory on Health Systems and Policies (www.euro.who.int) produces policy commentary and "health system reviews" on the health care systems of individual countries.

differ in both their populations and other health determinants. Evidence comparing public versus private health care provision within the same country has also been scarce. A recent working paper by Frakes, Gruber, and Justicz (2020) provides a rare quasi-experimental examination of this important question.[3] Studying military mothers who give birth in two different hospitals due to a move between deliveries, they find higher spending but lower rates of complications in private hospitals. In other recent work, event-study evidence suggests that the privatization or "corporatization" of health care providers in the US fee-for-service context leads to higher spending, worse outcomes, and worse access (Eliason et al. 2020; Andreyeva et al. 2022; Duggan et al. 2023).[4]

Second, an important literature has sought to measure the quality of care in the VA, which budgeted \$84 billion for medical care in 2020 (Department of Veterans Affairs 2020).[5] Following a well-known reorganization and investment in health IT in the mid-1990s (McCarthy and Blumenthal 2006), this literature has documented favorable VA quality, compared to care outside of the VA, in terms of process measures and health outcomes (e.g., Jha et al. 2003). The question of performance in the VA health care system has become particularly relevant in recent years, as the Department of Veterans Affairs considers ways to improve access to care for veterans and as Congress has sought to increase private health care delivery for veterans (Veterans Access, Choice, and Accountability Act of 2014; House Veterans' Affairs Committee 2018). So far, however, this literature has mainly compared outcomes of veterans receiving care in the VA system to those of nonveterans outside of the VA.

A third and extensive literature studies why, among developed countries, the US appears particularly inefficient in its health care, spending more as a percentage of GDP than any country, yet achieving poor outcomes relative to this spending (Garber and Skinner 2008). Experts have drawn attention to fragmentation in the US health care system, potentially increasing spending and worsening outcomes (Cebul et al. 2008; Cutler 2010; Agha, Frandsen, and Rebitzer 2019). Policymakers have responded by incentivizing the adoption of health IT and integrated care, but whether such policies improve health outcomes remains an open empirical question.[6] Our results are consistent with a productivity advantage (better outcomes at

[3] A related but distinct question of the impact of competition on government-provided health care is addressed in several important papers studying the British setting (e.g., Gaynor, Moreno-Serra, and Propper 2013; Bloom et al. 2015; Gaynor, Propper, and Seiler 2016). Within this strand of research, Cooper, Gibbons, and Skellern (2018) and Kelly and Stoye (2020) study the impact of private competition on public performance; both papers also document the selection of healthier patients to private hospitals.

[4] Eliason et al. (2020) find that acquisition of independent dialysis facilities by large chains leads to the facilities increasing doses of highly reimbursed drugs, replacing high-skill nurses with less-skilled technicians, and wait-listing fewer patients for kidney transplants. Patients under care of acquired facilities experience worsening mortality and hospitalizations. Andreyeva et al. (2022) find that acquisition of hospitals by large private systems increases hospital profitability through raising prices and intensifying the claimed intensity of procedures, while reducing the role of labor. Patients at these hospitals experience an increase in readmissions. Duggan et al. (2023) study 258 hospital privatizations and show evidence of greater profitability following privatization—stemming mostly from greater revenue per patient—but reduced access for Medicaid patients in the area.

[5] Spending continues to grow. The 2019 enacted budget allocated \$77 billion for VA medical care, and the 2021 proposed budget requests \$94 billion for medical care. In the last ten years, spending on medical care has nearly doubled (Department of Veterans Affairs 2020).

[6] A recent empirical literature documents modest reductions in spending and improvements in patient satisfaction among providers forming ACOs (McWilliams et al. 2014, 2016; Trombley et al. 2019). A mixed literature on health IT adoption has shown health improvements in some cases (e.g., Miller and Tucker 2011) but null results in general (e.g., Agha 2014). To our knowledge, our paper is the first to assess the complementarity between health IT and continuity of care.

lower cost) at the VA, the nation's largest integrated health care system. We find striking differences in reported services in VA versus non-VA settings: VA hospitals are much more likely to report utilization of low-cost services that improve coordination and continuity of care; non-VA hospitals are much more likely to report highly intense services. We also find suggestive evidence that government regulations to incentivize private hospitals to adopt health IT and integrate care may have improved outcomes among veterans with continuity of care at these hospitals.

The remainder of this paper proceeds as follows. Section I describes the setting and data. Section II presents our main analysis of the VA survival benefit. Section III discusses our survival analysis over time. Section IV presents evidence on mechanisms driving the VA survival benefit. Section V discusses policy implications and concludes.

I. Setting and Data

A. *US Health Care and the Veterans Health Administration*

The US health care system is marked by a high level of complexity, involving multiple private and public (federal, state, and local) parties. The US spends more on health care per capita than any other country—50 percent greater than the second-highest country, Norway—but has lower life expectancy than most other high-income countries (Rice et al. 2013). Compared to other high-income countries, the private sector plays a greater role in the US health care system.

Nonetheless, veterans in the United States have access to an important system of public provision: the Veterans Health Administration of the US Department of Veterans Affairs (VA). The VA provides health care for 9 million veteran enrollees, a number that has grown dramatically in recent decades (Chan, Duggan, and Guo 2021). The VA is the nation's largest integrated health care delivery system, including 170 medical centers and more than 1,000 outpatient sites of care, with a budget of $84 billion in 2020 for medical care (Department of Veterans Affairs 2020).

Key institutional features distinguish the VA from private providers in the United States. Like other systems of public provision, the VA has a well-defined patient population—enrolled veterans—while the link between private providers and their patients is complex and transient, driven by patient choice and insurance arrangements (Cebul et al. 2008; Agha, Frandsen, and Rebitzer 2019). The VA serves patients who are more disadvantaged (e.g., less likely to be employed, have lower income) than other patients (Agha et al. 2000). For most VA enrollees, care is essentially free in the VA; in contrast, the average Medicare beneficiary spends more than 40 percent of Social Security income on out-of-pocket health care costs (Cubanski et al. 2018; Hynes et al. 2021).[7] Local VA hospital systems receive financing according to the number of enrolled veterans and pay their physicians by salary (Wasserman et al. 2005); in contrast, private providers receive fee-for-service financing from an array of public and private insurers (Rice et al. 2013). This difference is consistent

[7] VA out-of-pocket payments depend on a veteran's priority group, income level, and disability ratings. VA cost-sharing rates can be found at https://www.va.gov/health-care/copay-rates/.

with payment arrangements across public and private delivery systems in other countries.[8]

In several ways, the VA's population-based contracting enables it to function as an integrated system, in contrast to most of the US health care system.[9] First, the VA directly employs all of its physicians and health care workers, while most physicians outside of the VA are independent of the hospitals at which they work and can affiliate with multiple hospitals. Second, care in the VA is integrated across clinical settings (e.g., inpatient, emergency department, outpatient) and across specialties of care. Third, in the mid-1990s, the VA implemented one of the first and most widely used electronic health record (EHR) systems in the United States—electronically connecting providers across locations and clinical settings. In comparison, only 1.5 percent of private US hospitals maintained a comprehensive EHR system prior to the Affordable Care Act (ACA) (Jha et al. 2009).[10]

B. *Comparing VA and Non-VA Care*

Over the past decade, lawmakers have enacted major reforms that allow veterans to receive VA-funded care at private facilities (Veterans Access, Choice, and Accountability Act of 2014; House Veterans' Affairs Committee 2018).[11] These reforms shift the VA's role to that of an *insurer* for veterans (similar to the role of Medicare for the elderly), with accompanying functions of authorizing care, processing claims, and detecting waste and fraud.

Related to these initiatives, the quality of care in the VA has been a long-standing subject of interest to policymakers and researchers. The health services literature has documented that the VA provides care of the same or higher quality than that of the private sector, as measured by a wide variety of process measures and health outcomes.[12] However, these comparisons are potentially confounded by differences, due to eligibility and self-selection, between the populations that utilize care in the VA and in non-VA facilities. Indeed, the vast majority of existing research has compared the care of veterans in the VA with the care of nonveterans in non-VA facilities.[13]

[8] Blank, Burau, and Kuhlmann (2017) provide a survey of health care systems across the world, describing countries by their degree of public provision (Figure 3.2) and their form of provider payment (Table 2.7). At the time of their survey, all countries with a high degree of public provision (Australia, Britain, New Zealand, Singapore, and Sweden) predominantly paid physicians by salary. All countries with a low degree of public provision (Japan, the Netherlands, Taiwan, and the US) paid physicians predominantly through fee-for-service contracts.

[9] The few notable cases of integrated delivery systems in the private sector (e.g., Kaiser Permanente, Intermountain Healthcare) also feature well-defined populations and financial integration with a large payer (Hendricks 1993; Bohmer, Edmondson, and Feldman 2013).

[10] While later federal policies sought to spur care coordination and health IT adoption in the private sector (Blumenthal 2010; Greaney 2011), most private hospitals that have adopted health IT still do not share records with each other (Holmgren, Patel, and Adler-Milstein 2017).

[11] There have been additional well-funded efforts to shift care further into the private sector (Kefe 2018; Rein et al. 2018; Shulkin 2018; Gordon 2019). According to an official recommendation to the congressionally established Commission on Care, some have even proposed that "if veteran choice dictates it over time, the long term goal of the transformation is the total transition to community care" (Blom 2016, p. 16).

[12] See Shekelle et al. (2010); Trivedi et al. (2011); and O'Hanlon et al. (2017) for systematic reviews. The literature includes dozens of studies on hundreds of quality of care process measures, as well as several studies on health outcomes.

[13] Two studies are noteworthy for having better identification. Nuti et al. (2016) compare outcomes for veterans in VA hospitals with outcomes for nonveterans in non-VA hospitals but restrict comparisons between VA and non-VA hospitals in the same metropolitan statistical areas. In an older study, Wright et al. (1999) look at 47,598

We use two key ideas to extend the literature on comparisons between VA and non-VA care. First, we focus on dually eligible veterans who are aged 65 and older. These veterans can receive care in the VA and at non-VA hospitals accepting Medicare (Hynes et al. 2007). A large and growing proportion of dually eligible veterans uses care in both VA and non-VA settings (Liu et al. 2018). Second, we build on the ambulance design strategy of Doyle et al. (2015) to sidestep concerns about the endogenous selection of where to obtain care. Specifically, we study veterans who arrive at a hospital via a 911-dispatched ambulance, comparing veterans from the *same zip code* who could have been transported by different ambulance companies with different propensities to transport patients to a VA hospital. Importantly, Doyle et al. (2015) document plausibly quasi-random variation in the assignment of patients to ambulance companies due to rotational arrangements, direct competition, and the locations of available ambulance units at the time of the 911 call. Ambulance companies also exhibit different tendencies to transport patients to various hospitals, based on their ownership, headquarters location, and other characteristics (Skura 2001). We further describe our quasi-experimental design and assess its assumptions in Section II.

C. *Data*

We use data from two main sources—Medicare claims and VA administrative data—for the universe of enrolled veterans in the VA from 2000 to 2014. We observe all Medicare claims for any dually enrolled veteran. These claims data include the beneficiary's zip code and demographic information (e.g., age, race, gender) as well as a record of medical services, each defined by an encounter date, Current Procedural Terminology (CPT) code(s), diagnostic (International Classification of Diseases, Ninth Revision, or ICD-9) codes, and provider identity. On the VA side, we have a complete record of clinical encounters in the electronic health record system that we transform into a corresponding set of encounter dates, CPT codes, ICD-9 codes, and provider identities.[14]

We begin by selecting ambulance ride events for dually eligible veterans, as recorded in the Medicare claims.[15] We restrict attention to "lights and sirens" emergency ambulance rides originating from 911 dispatch calls.[16] As in Doyle et al. (2015), we extract the date of the ambulance ride and the identity of the ambulance company, based on its tax identification number (TIN). We use the ambulance company identity to develop our instrumental variable for the propensity of the ambulance company to deliver patients to the VA or to non-VA hospitals. We also extract

dually eligible veterans with myocardial infarction. These studies find no difference or slightly better mortality outcomes in VA hospitals. Of note, a related literature suggests that veterans generally have poorer health than nonveterans (e.g., Agha et al. 2000).

[14] The VA system includes patient home address information. However, since we use Medicare claims for ambulance rides as the starting point for our sample, and the Medicare address information is updated regularly, we use Medicare claims records as our source for home location.

[15] VA policy is that patients with outside insurance should have ambulance services paid for by that insurance. In our dually eligible population, therefore, ambulance rides will be recorded in the Medicare claims.

[16] We select ambulance rides with Healthcare Common Procedure Coding System (HCPCS) codes A0322, A0328, A0330, A0362, A0368, A0370, A0427, A0429, A0433, or Q3019. We restrict to modifiers "RH," "SH," "NH," and "EH," corresponding to rides to a hospital from a residential location, a scene of an accident or acute event, a skilled nursing facility, and an extended care facility, respectively.

information on interventions provided by the ambulance (e.g., intravenous fluids, intubation), the level of care (advanced life support or basic life support), the pickup location (i.e., private residence, nursing home, skilled nursing facility, accident site), and the ambulance diagnosis (ICD-9) codes assigned by the ambulance personnel.

We then link these ambulance rides to emergency department (ED) visits at VA and non-VA hospitals.[17] Transport to the VA constitutes our treatment of interest. For each patient, we collect information on medical conditions and outpatient, ED, and inpatient utilization over the prior year, as recorded in the Medicare claims and VA records. We use the ICD codes for past medical conditions to identify 31 Elixhauser indices (Elixhauser et al. 1998) of comorbidities, noting the source of each condition (i.e., from visits to the VA, to non-VA facilities, or both). These comorbidities range from common conditions such as hypertension to rarer ones such as lymphoma.

Our primary outcome measure is mortality. We obtain information on the date of death from three sources: VA clinical records and Medicare claims, the Veterans Benefits Administration (VBA), and the Social Security Administration (SSA). The latter two sources are particularly reliable. They determine whether the veteran will receive payments from either the VBA or the SSA, and they draw on reports from family, funeral directors, post offices, financial institutions, other federal agencies, and state vital records agencies.

To construct our main analytical sample of 401,319 ambulance rides, we make the following restrictions (see online Appendix Table A.1). First, we remove patients who live in zip codes more than 20 miles away from either the nearest VA hospital or the nearest non-VA hospital. We also drop patients who traveled more than 50 miles from their zip code to the hospital. Second, we require that patients live in zip codes served by at least two ambulance companies with at least 20 rides in our data, at least 5 percent of rides transported to a VA hospital, and at least 5 percent transported to a non-VA hospital. Finally, for our baseline analysis of VA versus non-VA care, we drop veterans with no VA primary, ED, or inpatient care in the prior year, since ambulances transport fewer than 1 percent of these veterans to the VA.[18]

Table 1 describes the characteristics of the veterans and their emergency episodes at different steps in the creation of our main analytical sample. The 28-day mortality rate is stable across steps and relatively high, between 9.7 and 11.5 pp, reflecting the illness acuity of elderly veterans who arrive by a 911-dispatched ambulance. Similarly, the proportion of ambulance rides on weekend days is remarkably stable and close to two-sevenths, reflecting the unplanned nature of these health events (Card, Dobkin, and Maestas 2009). The major impact of our sample restrictions is to increase the share of rides going to a VA hospital. Notably, rides transporting veterans with both VA and non-VA ED visits in the prior year comprise about one-third of the rides transporting veterans with any ED visit in that year. In some steps, such as restricting to zip codes close to VA and non-VA hospitals, the sample becomes more concentrated in urban areas with shorter distances to nearby VA and non-VA

[17] We are able to link 90.2 percent of the ambulance rides with an ED visit in the VA or Medicare data.

[18] In a secondary analysis of continuity of care outside of the VA, in Section IV, we study an analogous sample of 1,414,217 ambulance rides of veterans who did not use VA care in the previous year and live in zip codes with at least 2 non-VA hospitals within 20 miles. Online Appendix Table A.14 describes the selection process for this sample.

Table 1—Characteristics of Baseline Sample

Restrictions	Sample characteristics				
	Dually eligible	Add zip × hospital	Add zip × ambulance	Add VA prior utilization	Add no ride in prior month
Male	0.899	0.883	0.863	0.962	0.963
Age	77.04	76.89	76.13	75.62	76.03
Black	0.111	0.163	0.187	0.200	0.194
Income	$21,724	$21,453	$20,874	$20,243	$20,905
Rural zip code	0.255	0.043	0.045	0.050	0.051
Residential source	0.610	0.600	0.652	0.685	0.705
Comorbidity count	6.53	6.69	6.44	6.54	6.14
Advanced life support	0.696	0.655	0.655	0.674	0.684
Weekend rate	0.272	0.269	0.270	0.270	0.269
Prior-year utilization					
VA ED only	0.048	0.082	0.130	0.277	0.294
Non-VA ED only	0.607	0.560	0.492	0.252	0.247
VA and non-VA ED	0.088	0.115	0.134	0.288	0.235
Ambulance rides	2.77	3.05	3.25	3.12	2.16
Outcomes					
28-day mortality	0.115	0.109	0.104	0.100	0.097
Present at VA	0.044	0.088	0.166	0.336	0.330
Observations					
Patients	2,862,557	1,118,302	365,163	188,299	188,299
Ambulance rides	8,828,997	3,465,588	1,051,093	491,193	401,319

Note: This table presents characteristics of observations remaining at each step of creating the baseline sample, detailed in online Appendix Table A.1.

hospitals. Black veterans also comprise a larger share of the sample. Patient characteristics otherwise remain stable across sample restriction steps.

Table 2 shows the characteristics and outcomes of patients in the main analytical sample by whether they arrive at the VA or a non-VA hospital. Veterans who arrive at the VA are more likely to be Black and have lower incomes on average. They have greater VA attachment and fewer recorded comorbidities, noting that comorbidities could be selectively recorded at non-VA hospitals for billing purposes. They have lower 28-day mortality, more hospital days and outpatient visits following the ambulance ride, and higher spending.

II. Benchmark Analysis

A. *Research Design*

Following Doyle et al. (2015), our empirical strategy relies on the assignment of ambulances to patients in emergencies and the role of ambulance companies in determining the hospital that provides emergency care to these patients. We assume (and attempt to verify empirically) that the assignment of a particular company is as good as random, conditional on zip code, ambulance characteristics, place of pickup, and patient use of VA and non-VA services in the past year.

TABLE 2—SAMPLE CHARACTERISTICS BY HOSPITAL DESTINATION

		Destination	
	Overall	VA	Non-VA
Male	0.963	0.960	0.964
Age	76.03	75.32	76.38
Black	0.194	0.254	0.164
Income	\$20,905	\$17,571	\$22,642
Rural zip code	0.051	0.057	0.048
Residential source	0.705	0.721	0.697
Comorbidity count	6.14	5.74	6.34
Prior VA ED visit only	0.294	0.563	0.161
Prior non-VA ED visit only	0.247	0.038	0.350
Prior VA and non-VA ED visit	0.235	0.237	0.234
Ambulance rides in prior year	2.16	2.23	2.12
Advanced life support	0.684	0.625	0.714
Weekend rate	0.269	0.265	0.271
Present at VA	0.330	1.000	0.000
28-day mortality	0.097	0.079	0.106
Admission	0.589	0.580	0.594
Hospital days	4.380	4.800	4.173
ED revisits	0.318	0.323	0.316
Outpatient visits	1.443	1.634	1.350
Spending	\$12,265	\$13,158	\$11,824
Number of patients	188,299	71,523	144,146
Number of ambulance rides	401,319	132,535	268,784

Notes: This table presents characteristics of observations in the baseline sample, detailed in online Appendix Table A.1. The first column of this table matches characteristics in the last column of Table 1. The second and third columns present characteristics in subsamples defined by whether the ride destination is the VA or a non-VA hospital.

Specifically, we condition on the origin zip code $\ell(i)$ of ambulance ride i so that we compare patients from the same zip code but transported by different ambulance companies. We also condition on whether the ambulance offers advanced life support (ALS) or basic life support (BLS) based on ambulance Healthcare Common Procedure Coding System (HCPCS) codes. We further condition on the ride's pickup site category (e.g., residential address, nursing home, scene of an accident), the day of the week, and month-year interactions (e.g., January 2010). Finally, we condition on coarse indicators of the patient's primary care, ED, and inpatient utilization at VA and non-VA facilities over the past year.[19] For simplicity, we refer to the joint set of controls for ambulance type, site of pickup, date of pickup, and prior utilization as "baseline controls" and denote them with $\mathbf{X}_i^0$. We describe the baseline controls in detail in online Appendix Table A.2.

Unlike Doyle et al. (2015), we do not include patient demographics, prior medical conditions, or ambulance diagnoses in the set of baseline controls. Instead, we

[19] We include these controls for two reasons. First, prior utilization may capture ambulance service areas within large zip codes. For example, riders within a large zip code who are closer to the VA may have greater VA prior utilization and may be served by a different set of ambulances. Second, these controls mirror our sample selection requirement that the veteran must have some VA utilization in the prior year (online Appendix Table A.1). We include more detailed measures of prior utilization in our "holdout" controls, described below. Below, we show that our results are qualitatively unchanged regardless of whether we include either set of measures of prior utilization.

"hold out" these variables—many of which are highly predictive of mortality—and use them to test whether there is balance across ambulance companies with differing propensities to send patients to the VA, conditional on $\left(\ell(i), \mathbf{X}_i^0\right)$.

To proceed more formally, let $D_i \in \{0,1\}$ represent an indicator for delivery to a VA hospital for ambulance ride i. We note that transfers between VA and non-VA facilities after the initial ED visit are rare in our sample and about equally likely for patients who go to VA and non-VA EDs.[20] Company $j(i) \in \mathcal{J}_{\ell(i)}$ provides ride i and is drawn from the set of companies $\mathcal{J}_\ell$ serving zip code ℓ.[21] Associated with each ride and company is a potential treatment indicator $D_i(j)$; thus, $D_i = D_i(j(i))$. Our main outcome is the 28-day mortality of the patient, denoted by $Y_i \in \{0,1\}$. The associated potential outcomes $Y_i(d)$, $d \in \{0,1\}$, depend on whether the patient was transported to a VA hospital $(d = 1)$ or not $(d = 0)$, with $Y_i = Y_i(D_i)$.

Under the assumption that different ambulance companies have systematically different tendencies to transport patients to the VA, and that the assignment of $j(i)$ is as good as random, conditional on $\left(\ell(i), \mathbf{X}_i^0\right)$, we can use the identity of the ambulance company to construct a valid instrumental variable for D_i. We consider the following conditions for IV validity (Imbens and Angrist 1994).

CONDITION 1 (IV Validity): *For a random sample of ambulance rides i provided by ambulance companies j, the following conditions hold:*

(*i*) Relevance: $E\left[D_i(j) \mid \ell(i), \mathbf{X}_i^0\right]$ *is a nontrivial function of* $j \in \mathcal{J}_{\ell(i)}$.

(*ii*) Independence and Exclusion: *The vector of potential outcomes*, $\left(Y_i(0), Y_i(1), D_i(j)\right)$, *is independent of the assigned ambulance company*, $j(i) \in \mathcal{J}_{\ell(i)}$, *conditional on* $\left(\ell(i), \mathbf{X}_i^0\right)$.

(*iii*) Monotonicity: *Conditional on* $\left(\ell(i), \mathbf{X}_i^0\right)$, *for any j and j′*, $D_i(j) \geq D_i(j')$ *for all i, or* $D_i(j) \leq D_i(j')$ *for all i.*

We will discuss each of these conditions and present evidence supporting them in Section IIB.

Our research design adopts the same structure as studies that exploit the random assignment of judges (who vary in terms of leniency) for the purpose of identifying the impact of some court-determined treatment (e.g., Kling 2006; Dahl, Kostol, and Mogstad 2014). As is standard in this judges-design literature, to deal with finite samples, we construct a leave-out (or jackknife) instrumental variable that reflects the propensity of the ambulance company $j(i)$ assigned to ride i to transport *other* patients to the VA. We compute this as the average fraction of other patients

[20] Of the 132,535 rides that go to the VA, 828 (0.6 percent) have a non-VA hospital ED visit by the subsequent day. Of the 268,784 rides to a non-VA hospital, 2,191 (0.8 percent) have an ED visit at a VA hospital by the next day. Of 79,684 VA admissions, 418 (0.5 percent) were transferred to a non-VA hospital within 7 days of the ambulance ride. Of 157,682 non-VA admissions, 1,774 (1.1 percent) were transferred to a VA facility within 7 days.

[21] We define an "ambulance company" by the interaction between the tax identification number (TIN) and the health referral region (HRR) associated with the ride. This definition accounts for a few large corporations with a single TIN that serve multiple regions.

picked up by company $j(i)$ and transported to the VA. Specifically, for ambulance ride i transporting patient $k(i)$, we define the leave-out probability Z_i of transport to the VA:

$$Z_i = \frac{1}{K_{j(i)} - 1} \sum_{i' \in \mathcal{I}_{j(i)}} \frac{\mathbf{1}\{k(i') \neq k(i)\} D_{i'}}{N_{k(i'),j(i)}}, \tag{1}$$

where K_j is the total number of patients transported by company j, $N_{k,j}$ is the total number of rides taken by patient k with company j, and $\mathcal{I}_j$ is the set of rides transported by ambulance company j. We estimate Z_i using the sample of dually eligible veteran ambulance rides (column 1 of Table 1).

Under Condition 1, an IV estimate based on Z_i, conditioning on $(\ell(i), \mathbf{X}_i^0)$, recovers a local average treatment effect (LATE) of the VA on mortality among compliers. For comparison, we also consider the observational "treatment effect" of going to the VA on mortality of patients taken to a hospital by a 911-dispatched ambulance, controlling for $(\ell(i), \mathbf{X}_i^0)$:

$$Y_i = \beta D_i + \mathbf{X}_i^0 \delta_0 + \zeta_{0,\ell(i)} + \varepsilon_{0,i}, \tag{2}$$

where $\zeta_{0,\ell}$ represents an unrestricted fixed effect for rides originating in zip code ℓ. As in Doyle et al. (2015), zip code fixed effects imply that both our observational and quasi-experimental concepts of the VA treatment effect involve comparisons between veterans who live in the same zip code. Estimating equation (2) by OLS yields $\hat{\beta}_{OLS}$, while instrumenting D_i with Z_i yields $\hat{\beta}_{IV}$.

The compliers in this quasi-experiment are rides with dually eligible veterans who could be swayed to a VA or non-VA hospital, depending on the identity of the ambulance company. Because VA hospitals typically only treat veterans, these compliers are patients who may state to an ambulance (but possibly not to all ambulances) that they are veterans and would be open to care at the VA. By definition, the compliers exclude veterans who would insist on being taken to a VA hospital or a hospital outside of the VA. Nonetheless, in the following subsection, we estimate a sizable share of compliers, consistent with the high percentage of veterans in the baseline sample with both VA and non-VA ED visits in the prior year (Table 1). We also estimate and report complier characteristics in Section IIA. As expected, compliers are more likely to have had previous ED visits at the VA; almost three in ten compliers have had ED visits in both VA and non-VA hospitals in the prior year. Differences in veteran characteristics between compliers and those who select a VA or a non-VA hospital suggest that D_i is not quasi-experimentally assigned and that OLS may be biased. We discuss support for Condition 1(ii)—that Z_i is quasi-experimentally assigned—in the next subsection.

The gap between $\hat{\beta}_{OLS}$ and $\hat{\beta}_{IV}$ may stem from differences in the potential outcomes between never-takers (i.e., patients who go to a non-VA facility regardless of the ambulance company), always-takers (i.e., patients who go to the VA regardless of the ambulance company), and compliers. In the setting of a VA advantage, if the always-takers are sicker than the never-takers, then $\hat{\beta}_{OLS}$ may be smaller than $\hat{\beta}_{IV}$. We explore this gap more directly in Section III.

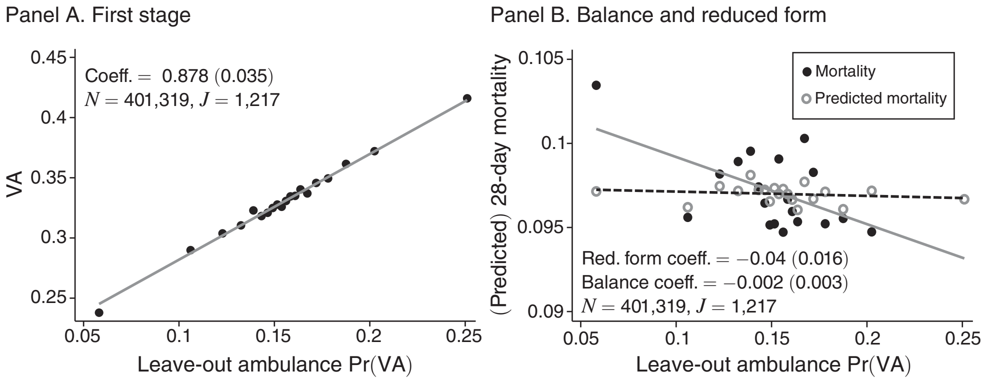

Figure 1. First Stage, Balance, and Reduced Form

Notes: Panel A shows a binned scatterplot of arrival at a VA hospital on the y-axis against the ambulance leave-out propensity to arrive at a VA hospital on the x-axis. The figure is a graphical representation of the first-stage regression in equation (3). Panel B shows binned scatterplots of 28-day mortality and predicted 28-day mortality on the y-axis against the ambulance leave-out propensity to arrive at a VA hospital on the x-axis. Mortality bin means are shown in solid circles, while predicted mortality bin means are shown in hollow circles. The figure represents the reduced-form regression in equation (4) and the corresponding balance regression replacing mortality with predicted mortality. The sample includes 401,319 ambulance rides and 1,217 combinations of ambulance company identifiers and Dartmouth Atlas Hospital Referral Regions (HRRs). The sample selection is given in online Appendix Table A.1. Baseline controls are detailed in online Appendix Table A.2 and include patient zip code dummies, ALS/BLS dummies, source of the ambulance ride, time categories, and patient prior utilization.

B. *First Stage, Balance, and Reduced Form*

We begin our empirical analysis by demonstrating instrument relevance, Condition 1(i), with the following first-stage regression:

$$D_i \;=\; \pi_1 Z_i + \mathbf{X}_i^0 \delta_1 + \zeta_{1,\ell(i)} + \varepsilon_{1,i}. \tag{3}$$

The coefficient π_1 reflects the impact of ambulance company preferences on the probability that the ride goes to the VA, conditional on our baseline controls for ambulance type, pickup site, zip code, date categories, and veteran prior utilization. Panel A of Figure 1 shows a binned scatterplot of residualized D_i on the y-axis and residualized Z_i on the x-axis and reports $\hat{\pi}_1 \;=\; 0.882$ (SE 0.034). The first-stage relationship between D_i and Z_i is very predictive and close to linear.

To assess independence, Condition 1(ii), we test whether Z_i is correlated with patient characteristics that predict mortality. Specifically, we construct an estimate of predicted mortality $\hat{Y}_i$ using "hold-out" patient characteristics of patient demographics and 31 Elixhauser indices for prior medical conditions.[22] We then fit models for $\hat{Y}_i$ based on the same right-hand-side specification as in equation (3).

[22] Patient demographics include age, gender, and race and ethnicity. We capture age by 2-year age bins from 65 years to 100 years. We capture race and ethnicity by three dummies for White, Black, and Hispanic; the omitted category is Asian/other. We use the 31 Elixhauser indices described in Elixhauser et al. (1998), interacting each index with the source of the comorbidity record. There are three possible sources: VA only, Medicare claims only, and VA and Medicare claims. This results in $3 \times 31 \;=\; 93$ dummies. Online Appendix Table A.3 further describes holdout patient characteristics.

Panel B of Figure 1 shows (with hollow dots) no relationship between $\hat{Y}_i$ and Z_i, controlling for $\left(\ell(i), \mathbf{X}_i^0\right)$.[23] In contrast, the same panel shows (with solid dots) that the reduced-form relationship between actual mortality, Y_i, and Z_i is significantly negative, under the same controls. Specifically, for the reduced-form relationship,

$$Y_i = \pi_2 Z_i + \mathbf{X}_i^0 \delta_2 + \zeta_{2,\ell(i)} + \varepsilon_{2,i}. \tag{4}$$

We find $\hat{\pi}_2 = -0.040$ (SE 0.016). This suggests that quasi-random assignment to an ambulance company more likely to transport to the VA results in an intention-to-treat reduction in mortality.

Under independence, we may quantify the share of compliers in our sample. As shown in panel A of Figure 1, 24 percent of rides assigned values of Z_i in the lowest vigintile still go to the VA. We may consider veterans in these rides as "always-takers." On the other hand, 58 percent of rides that are assigned values of Z_i in the highest vigintile still go to a non-VA hospital. We may consider veterans in these rides as "never-takers." The remaining share of rides, or 18 percent, characterizes the share of compliers.[24]

The exclusion assumption in Condition 1(ii) asserts that ambulance companies do not affect outcomes other than through their effect on whether a patient arrives at a VA or non-VA hospital. Our notation also implicitly assumes that each complier has a well-defined non-VA hospital that is stable across ambulance companies. In online Appendix A, we discuss weaker versions of the exclusion condition and show the robustness of our results to potential violations of it.[25] Specifically, we assess and find no evidence of any correlation between Z_i and ambulance treatments captured in summary charges or between Z_i and ambulance propensities to deliver patients to different non-VA hospitals. We also exploit the mortality outcomes of patients with no prior VA utilization. These patients are dropped from our main analytic sample because they have almost no chance of going to the VA, but they ride with the same ambulance companies as patients in our sample. Controlling for an ambulance's mortality outcomes among these out-of-sample patients may mimic controlling for exclusion violations. In all of these analyses, we find that our main IV estimate below is qualitatively unchanged, suggesting that our results are robust to violations of exclusion.

To assess the monotonicity assumption given by Condition 1(iii), we follow the standard practice in the judges-design literature to show that the first-stage

[23] In online Appendix Figure A.1, we present a simulation exercise that suggests that we have the power to reject a data-generating process in which more than 2–3 percent of patients are perfectly sorted by their predicted mortality to ambulances by the ambulance's propensity to transport to the VA.

[24] Interestingly, this share of compliers appears roughly similar to that in Doyle et al. (2015). Characterizing compliers in Doyle et al. (2015) is less straightforward because the treatment of hospital spending is a continuous variable. Nonetheless, if we characterize a treatment as a one standard deviation increase in spending (0.2 log points), this is slightly more than the difference between the first and fourth quartile of their instrument (Table 1 in Doyle et al. 2015). We could then interpret their first-stage coefficient of 0.17 (Table 2 in Doyle et al. 2015) as a complier share. While we note that compliers in our setting need to reveal that they are a veteran (to at least one ambulance), which may reduce the share of compliers, our sample of veterans is substantially more disadvantaged, which may imply a *higher* complier share (Card, Fenizia, and Silver 2023).

[25] Following Kolesar et al. (2015), the analyses in online Appendix A correspond to the weaker assumption that our instrument is uncorrelated with other ambulance-specific treatments impacting our outcome. Specifically, under this weaker version of exclusion, we require that ambulance companies with higher values of $E\left[D_i(j) \,\middle|\, \ell(i), \mathbf{X}_i^0\right]$ do not also systematically affect mortality by (i) observed treatments during the ambulance ride, (ii) delivery to higher- or lower-quality non-VA alternatives, or (iii) any other channel.

Table 3—Effect of VA Hospitals on Mortality

	Dependent variable: 28-day mortality				
	(1)	(2)	(3)	(4)	(5)
Panel A. OLS					
VA hospital	−0.024	−0.023	−0.026	−0.021	−0.021
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Outcome mean	0.097	0.097	0.097	0.097	0.097
Observations	401,319	401,319	401,319	401,319	401,319
Panel B. IV					
First stage	0.878	0.853	0.839	0.837	0.862
	(0.035)	(0.034)	(0.034)	(0.034)	(0.043)
IV estimate	−0.045	−0.037	−0.035	−0.038	−0.051
	(0.018)	(0.018)	(0.018)	(0.017)	(0.023)
Outcome mean	0.097	0.097	0.097	0.097	0.097
Observations	401,319	401,319	401,319	401,319	401,319
Controls					
Patient background	No	Yes	Yes	Yes	Yes
Prior diagnoses	No	No	Yes	Yes	Yes
Ambulance diagnoses	No	No	No	Yes	Yes
Co-riders characteristics	No	No	No	No	Yes

Notes: This table shows OLS and IV estimates of the effect of VA hospitals on 28-day mortality. Panel A gives OLS estimates, $\hat{\beta}_{OLS}$, for β in equation (2). Panel B gives IV estimates, $\hat{\beta}_{IV}$, as well as the first-stage coefficient, $\hat{\pi}_1$ in equation (3), with respect to the leave-out probability of the assigned ambulance company to transport patients to the VA. Baseline controls in all specifications are described in online Appendix Table A.2 and include patient zip code dummies, ALS/BLS dummies, source of the ambulance ride, time categories, and patient prior utilization. Patient background controls include demographics, socioeconomic status, combat history, eligibility for benefits, and counts of prior utilization. All additional controls are described in further detail in online Appendix Table A.3. The estimation sample is described in online Appendix Table A.1.

relationship between D_i and Z_i remains positive for subgroups of patients defined by different observable characteristics (e.g., Arnold, Dobbie, and Yang 2018; Bhuller et al. 2020). We detail these analyses in online Appendix B. Section III presents a stronger test of monotonicity (and IV validity) based on *potential outcomes*. Following the reasoning in Kitagawa (2015), this test amounts to showing a positive density for the potential outcome of death in a given week among compliers.

C. *Mortality Effect*

With this background, we now move to our main results on patient mortality. In Table 3, we show both OLS and IV estimation results for equation (2). Panel A of the table shows $\hat{\beta}_{OLS}$ from equation (2), while panel B shows $\hat{\beta}_{IV} = \hat{\pi}_2/\hat{\pi}_1$ from the first-stage and reduced-form regressions in equations (3) and (4). Column 1 shows our baseline specification, controlling for zip code and the variables in $\mathbf{X}_i^0$. The OLS estimate is $\hat{\beta}_{OLS} = -0.024$ (SE 0.001), while the IV estimate is $\hat{\beta}_{IV} = -0.045$ (SE 0.018).[26] Relative to the mean 28-day mortality of 9.7 pp, both estimates imply a sizable reduction in mortality for compliers taken to the VA.

[26] Online Appendix Figure A.2 shows the IV estimate visually by plotting the predicted first-stage probability of treatment from equation (3) on the x-axis and the predicted reduced-form effect on mortality from equation (4) on the y-axis. The slope of this visual IV relationship corresponds to $\hat{\beta}_{IV} = -0.045$.

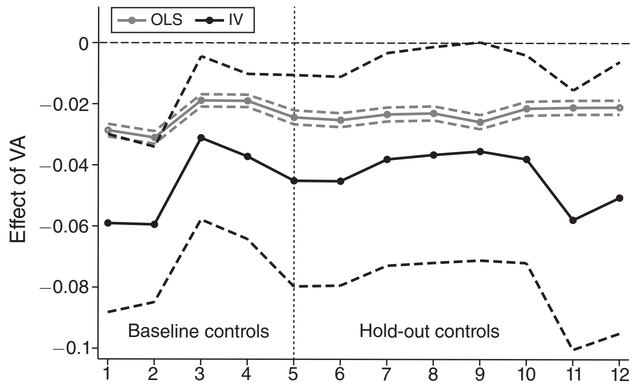

FIGURE 2. OLS AND IV SPECIFICATIONS

Notes: This figure shows OLS and IV estimates of the effect of the VA on 28-day mortality, represented in equation (2) as β, with progressive sets of controls. Numbered incremental controls correspond to categories or subcategories of variables presented in order in online Appendix Tables A.2 and A.3. Control sets are as follows: (1) zip code; (2) pickup source; (3) ambulance service; (4) time categories; (5) prior utilization; (6) demographics; (7) socioeconomic status, combat history, and eligibility; (8) extended prior utilization; (9) prior diagnoses; (10) three-digit ambulance diagnosis codes; (11) co-rider baseline controls; and (12) co-rider hold-out controls. Estimates are shown along solid lines, while 95 percent confidence intervals are shown in dashed lines. All specifications use the baseline sample, given in online Appendix Table A.1.

The other columns in Table 3 show OLS and IV estimates as we include additional controls to the models: (i) patient demographics (age, race, gender), socioeconomic status, combat history, benefit eligibility, and finer measures of prior utilization; (ii) ambulance diagnostic (ICD-9) codes; (iii) Elixhauser comorbidity indicators; and (iv) ambulance and co-rider controls, which are all described in online Appendix Table A.3.[27]

Reassuringly, both $\hat{\beta}_{OLS}$ and $\hat{\beta}_{IV}$ remain stable as we add additional controls. Figure 2 illustrates this stability as we add controls in a more granular fashion; online Appendix Figure A.3 shows the stability of the IV estimates as we permute the order in which we add controls. The stability of both the OLS and IV estimates suggests a lack of selection based on a wide range of observable patient and co-rider characteristics. If anything, the inclusion of co-rider controls (shown as control sets (11) and (12) in Figure 2) slightly *increases* the magnitude of $\hat{\beta}_{IV}$, though the difference is not statistically significant. Under the reasoning of Altonji, Elder, and Taber (2005), this stability suggests limited scope for selection on unobservable characteristics that predict potential 28-day mortality.

Further, our IV estimates are larger than our OLS estimates. While a Hausman test for equality of the two estimates with baseline controls has a t-statistic of only 1.17, we show in the next section a dynamic pattern of IV and OLS estimates, using weekly data over the year following the initial ambulance ride, which points more definitively to systematic differences. Using this approach, we can infer with high

[27] Following the reasoning in Altonji and Mansfield (2018), the co-rider controls can capture patient selection at the ambulance company level beyond the observable characteristics of the index patient by using characteristics of *other* rides and patients under the same ambulance company. Specifically, these controls address the concern that sicker patients may be allocated to ambulance companies that systematically differ in their propensity to transport patients to the VA.

confidence that the causal VA advantage is larger than the (precisely estimated) OLS effect. The difference between OLS and IV estimates could arise for two reasons: (i) never-takers are healthier than always-takers (i.e., selection runs *counter* to treatment effects on mortality), or (ii) the LATE is larger than the unconditional average treatment effect (ATE). We investigate these possibilities in the next section.

Our IV estimate of the VA effect on mortality is large but within the range of prior mortality effects in acute settings.[28] The studies closest to ours are Doyle et al. (2015) and Hull (2020), both of which estimate mortality effects among hospitals included in the Medicare data using the ambulance design. Doyle et al. (2015) find that hospitals with higher log spending by 1 standard deviation decrease 1-year mortality by 3.7 pp, about 10 percent of average 1-year mortality in their sample. Hull (2020) studies the distribution of hospital effects on 30-day mortality and finds effects similar in absolute size but larger when stated as a percent of the sample 30-day mortality.[29] Woodworth (2020) finds that a 10 percent reduction in ED crowding reduces 30-day mortality by 24 percent but has no significant effect on 1-year mortality. Doyle (2011) finds that local areas in Florida with higher spending by 2 standard deviations reduce in-hospital mortality by 26 percent. Card, Dobkin, and Maestas (2009) find that Medicare eligibility reduces 7-day mortality by 14 to 20 percent, 28-day mortality by 7 to 9 percent, and 1-year mortality by 2 to 4 percent.

A difficulty comparing mortality effects across the literature is reconciling differences in time frames for measuring mortality. In Section III, we perform a survival analysis that provides insight into the evolution of mortality following the ambulance ride. We show that mortality effects are concentrated in the first week. When we look at one-year mortality, the VA effect is slightly smaller than the Doyle et al. (2015) effect of higher-spending hospitals cited above.

III. Survival Analysis

In this section, we develop and apply a survival analysis framework to understand the dynamics of potential survival outcomes following the ambulance ride. We use this framework to draw several insights. First, we determine the time course of VA effects on mortality. Second, we use this framework to extend our validation of Condition 1, beyond our standard analysis in Section IIB. Third, we investigate the implications of heterogeneity in mortality risks between compliers and noncompliers of our ambulance quasi-experiment.

[28] Other papers have demonstrated significant mortality impacts outside of the hospital setting. Deryugina and Molitor (2020) show that moving to a region in the tenth versus the ninetieth percentile of local mortality corresponds to a 0.76 pp reduction in mortality, or 14 percent relative to the median of 5.3 pp Finkelstein, Gentzkow, and Williams (2021) show similarly large place-based impacts on mortality. Abaluck et al. (2021) show that reassignment from Medicare Advantage plans in the top ventile of mortality would reduce 1-year mortality by 27 percent. Other notable papers demonstrate mortality impacts of gaining Medicaid or health insurance more broadly (Sommers, Gawande, and Baicker 2017; Goldin, Lurie, and McCubbin 2021; Miller, Johnson, and Wherry 2021).

[29] Hull (2020) finds that, compared to admission to a random hospital, admission to hospitals with the highest-quality posterior estimates in a local market would reduce mortality by 3.6 to 4.7 pp, about 20 to 28 percent of the sample mean.

A. *Approach*

Consider a set of potential survival outcomes $S_i(t;d) \in \{0,1\}$ representing indicators for whether the patient in ride i would be alive in week $t \in \{1,\ldots,52\}$ following the ambulance ride, if taken to the VA $(d = 1)$ or a non-VA hospital $(d = 0)$.[30] By definition, if $S_i(t;d) < S_i(t-1;d)$, then the patient in ambulance ride i would die in the t^{th} week following the ambulance ride if exposed to treatment d. Of course, potential survival outcomes must weakly decrease over time, i.e., $S_i(t;d) \leq S_i(t-1;d)$ for all i, d, and t.

As with mortality outcomes, for each ambulance ride i, we can only observe the set of survival outcomes corresponding to $d = D_i$: $S_i(t) = D_i S_i(t;1) + (1-D_i)S_i(t;0)$. However, appealing to Abadie (2002), we can recover the expected survival outcomes for the set of compliers $\mathcal{C}$ whose hospital choice depends on which ambulance company picks them up. [31]

Given the potential survival outcomes, we can then estimate potential hazard rates for mortality under either VA or non-VA assignment:

$$h_{IV}(t;d) \equiv E\left[1 - S_i(t+1;d) \,\middle|\, S_i(t;d) = 1, i \in \mathcal{C}\right] = \frac{s_{IV}(t;d) - s_{IV}(t+1;d)}{s_{IV}(t;d)}, \tag{5}$$

for $d \in \{0,1\}$ and $t \in \{1,\ldots,52\}$, corresponding to weekly mortality hazard rates up to one year after the initial ambulance ride. Under Condition 1, differences between $\{h_{IV}(t;1)\}_t$ and $\{h_{IV}(t;0)\}_t$ can be interpreted as the causal effect of VA assignment, among compliers, on the set of mortality hazard rates.[32]

As in Section II, we also calculate risk-adjusted OLS survival functions and mortality hazard rates, conditional on D_i. We estimate $s_{OLS}(t;d) \equiv E\left[S_i(t;d) \,\middle|\, D_i = d\right] = E\left[S_i(t) \,\middle|\, D_i = d\right]$ by OLS, replacing the outcome variable in equation (2) with $S_i(t)D_i$ for $s_{OLS}(t;1)$ and with $S_i(t)(D_i - 1)$ for $s_{OLS}(t;0)$. Our OLS estimand of the VA effect on 28-day mortality, β_{OLS}, is similarly equal to $s_{OLS}(4;1) - s_{OLS}(4;0)$. Corresponding mortality hazard rates can also be calculated based on observed risk-adjusted survival:

$$h_{OLS}(t;d) \equiv E\left[1 - S_i(t+1;d) \,\middle|\, S_i(t;d), D_i = d\right] = \frac{s_{OLS}(t;d) - s_{OLS}(t+1;d)}{s_{OLS}(t;d)}. \tag{6}$$

[30] We adopt the convention that a mortality event within the first seven days occurs in week 1. Thus, a mortality event within 28 days occurs by the end of week 4.

[31] In particular, under Condition 1, we can estimate $s_{IV}(t;1) \equiv E\left[S_i(t;1) \,\middle|\, i \in \mathcal{C}\right]$ by two-stage least squares using the first-stage equation (3) and a reduced-form equation similar to equation (4) but with dependent variable $S_i(t)D_i$. Similarly, we can estimate $s_{IV}(t;0) \equiv E\left[S_i(t;0) \,\middle|\, i \in \mathcal{C}\right]$ using the same first-stage model but replacing the reduced-form outcome variable in equation (4) with $S_i(t)(D_i - 1)$. Note that, by construction, the IV estimand of the VA treatment effect on 28-day mortality in Section II satisfies $\beta_{IV} = s_{IV}(4;1) - s_{IV}(4;0)$.

[32] We emphasize that any gap between $h_{IV}(t,1)$ and $h_{IV}(t,0)$ at a later time horizon (e.g., $t = 12$) could arise because treatments at the VA affected the population of compliers who survive to week $t-1$ and are therefore at risk of death in week t, or because of a treatment effect on the week t hazard, holding the population fixed.

Compared to the potential survival functions and mortality hazards, the OLS analogs also incorporate outcomes for the always-takers and never-takers whose choice of hospital is unaffected by the specific ambulance company that picked them up. Specifically, $s_{OLS}(t;1)$ and $h_{OLS}(t;1)$ reflect survival outcomes for a combination of always-takers and compliers, while $s_{OLS}(t;0)$ and $h_{OLS}(t;0)$ reflect survival outcomes for a combination of never-takers and compliers.

B. *Time Course of Mortality Effects*

Since we examine potential survival outcomes one year after an ambulance ride, we restrict the analysis in this section to ambulance rides of patients with no ride in the prior year.[33] Figure 3 shows the estimated potential survival curves and potential hazard rates in weeks 0 to 52 for compliers assigned to the VA and those assigned to a non-VA hospital. The potential survival curves, shown in panel A, reveal a high risk of mortality among compliers. Mortality at 28 days among compliers assigned to a non-VA hospital is greater than the sample mean of 9.7 pp, and cumulative mortality at 1 year is approximately 30 pp. However, despite the substantial mortality risk over the subsequent year, the gap in survival between VA- and non-VA-assigned compliers (i.e., the mortality treatment effect) is fully realized at 28 days and remains stable for the rest of the year.

In panel B, we examine the implied hazard rates and show that the differences in mortality are concentrated in the first week following the ambulance ride.[34] Thereafter, though the hazard rates for both VA- and non-VA-assigned compliers remain relatively high, they are indistinguishable from each other. This similarity suggests that the VA advantage results entirely from events within the first week following the ambulance ride.

The potential hazard profiles in Figure 3 suggest that mortality risks for the compliers comprise two separate risks: (i) a relatively high short-term risk component that the VA reduces and (ii) a relatively stable long-term risk component that remains the same for compliers regardless of whether they are assigned to a VA or non-VA hospital. If the latter risk reflects underlying patient health and is independent of the risk that led to the ambulance call, then we would expect the long-run weekly mortality rate (e.g., after three months) to be the same for veterans quasi-randomly assigned to VA and non-VA hospitals. We formalize this as a test in Section IIIC.

The potential hazard rates allow us to assess whether excess mortality at non-VA hospitals involves "harvesting," or mortality displacement, in which deaths for

[33] This restriction attributes survival for a given patient in a given week to the "upstream" ambulance ride, rather than attributing the survival event to both upstream and downstream ambulance rides. This changes (decreases) the sample in online Appendix Table A.1 to 254,782 rides and 188,299 patients. In online Appendix Figure A.13, we show that this restriction (or any other restriction on prior rides) does not lead to qualitative differences in our estimated OLS or IV treatment effects on mortality over time. Regardless of the number of days within which we require no prior ride, the IV estimates are larger than 4 pp at 28 days and are stable within the year following the ambulance ride. The OLS estimates are between 2.0 and 2.5 pp at 28 days and essentially disappear by 1 year after the ambulance ride. We evaluate the implications of the long-term difference between IV and OLS treatment effects in Section IIID.

[34] In fact, most of this impact on mortality occurs on the first day. This pattern of survival benefits following an acute event is similar to the one demonstrated by Card, Dobkin, and Maestas (2007). Figure 10 of that paper shows that, in absolute terms, the survival benefit is almost fully present and highly significant by the first day of hospitalization.

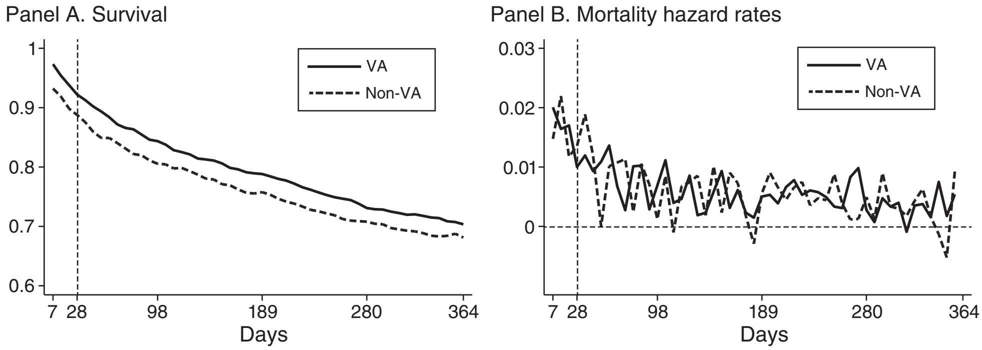

FIGURE 3. COMPLIER POTENTIAL OUTCOMES

Notes: This figure shows potential outcomes for ambulance compliers who arrive at a VA hospital and those who arrive at a non-VA hospital. Panel A shows survival outcomes as a function of days from the ambulance ride. "Days" indicate one-week intervals from the ambulance ride. Denote $S_i(t;d) \in \{0,1\}$ as an indicator for whether patient i survives up to time t after the ambulance ride, depending on whether the patient arrives at the VA ($d = 1$) or a non-VA hospital ($d = 0$). Observed survival is $S_i(t) = D_i S_i(t;1) + (1 - D_i) S_i(t;0)$. We estimate complier VA survival, or $E[S_i(t;1) \mid i \in \mathcal{C}]$, by an IV regression with a dependent variable of $S_i(t) D_i$, the endogenous VA treatment D_i, and the same first-stage and reduced-form design matrix implied by equations (3) and (4). We estimate complier non-VA survival, or $E[S_i(t;0) \mid i \in \mathcal{C}]$, by a similar IV regression with a dependent variable of $S_i(t)(D_i - 1)$. All regressions use a sample of ambulance rides with no prior ride in the last year and the same baseline controls as described in Figure 1. Panel B presents implied weekly mortality hazard rates, as given by equation (5).

patients at the VA are simply delayed (Schwartz 2000; Honore and Lleras-Muney 2006). Under this hypothesis, survival gains from VA care observed at 28 days are temporary and will fade in the long term. Such mortality displacement would imply that the hazard of dying *increases* among VA-assigned compliers relative to non-VA-assigned compliers after a time. We find no evidence of this in the potential hazard rates in panel B of Figure 3. In online Appendix A.2, we formally test that $h_{IV}(t;1) \leq h_{IV}(t;0)$ for all t and cannot reject this null hypothesis of no harvesting.[35] This suggests that the VA *prevents* rather than *displaces* deaths, leading to a persistent survival benefit, as shown by the stable gap between potential survival curves in panel A of Figure 3.

Finally, we can use the results in this section to put the magnitude of the VA advantage in context with the results in Doyle et al. (2015). Recall that Doyle et al. (2015) find that hospitals with 1 standard deviation higher log spending decrease 1-year mortality by 3.7 pp, or 10 percent of average 1-year mortality in their sample. In the sample of ambulance rides with no prior ride within a year, our estimates imply that the VA reduces mortality by 2.2 pp, or 7.7 percent of average 1-year mortality in our sample. The lower relative impact on mortality at 1 year, compared to the benchmark impact on 28-day mortality, is a consequence of much higher cumulative mortality at 1 year and the concentrated impact on mortality in the first week following the ambulance ride.

[35] Our test builds on Wolak's (1987) suggestion to form a test statistic based on a quadratic form that represents deviations of the data from the predictions of a constrained model that imposes the inequality restrictions. We use a simple bootstrap procedure to derive critical values of the test.

C. *Extended Test of IV Validity*

We can also use the estimated potential survival outcomes to test the validity of our IV strategy based on ambulance assignment. Under Condition 1, the density of any characteristic, including characteristics defined by potential outcomes, must be positive among compliers of the quasi-experiment (Balke and Pearl 1997; Imbens and Rubin 1997):

$$\Pr\left(X_i = x, Y_i = y \mid i \in \mathcal{C}\right) \geq 0, \tag{7}$$

for all possible characteristics $x \in \mathcal{X}$ and all possible potential outcomes $y \in \mathcal{Y}$. Kitagawa (2015) proposes a formal test of this implication, and Chan, Gentzkow, and Yu (2022) show that applying this test to *potential outcomes* can provide a stronger test of the conditions for IV validity, particularly the monotonicity assumption in Condition 1(iii).[36]

In our setting, we partition potential survival outcomes into weeks of potential mortality for each of the 52 weeks following the ambulance ride, for both VA- and non-VA-assigned compliers. Since survival can only decrease over time, the potential mortality hazard rates for any week must be positive; i.e., $h_{IV}(t;d) \geq 0$ for all $t \in \{0,\ldots,51\}$, $d \in \{0,1\}$. This prediction may be violated if patients' potential mortality in some week t is correlated with their assigned ambulance's propensity to go to the VA (a violation of independence). It may also be violated if there exist "defiers" (i.e., patients who are *less likely* to go to the VA when assigned to ambulances that transport more often to the VA overall) with a higher risk of death in some week t (a violation in monotonicity). In online Appendix A.2, we formally test the joint inequality constraint that $h_{IV}(t;d) \geq 0$ for all $t \in \{0,\ldots,51\}$, $d \in \{0,1\}$, and cannot reject this null hypothesis, with a bootstrap-based p-value of 1.00.

If the short-term and longer-term mortality risks facing veterans are independent (as is typically assumed in a competing risks model) and treatment at the VA only affects the short-term risk component, then Condition 1 also implies that $h_{IV}(t;1) = h_{IV}(t;0)$ for $t \geq \bar{t}$, for some $\bar{t}$ after the acute ambulance episode. Specifically, if ambulance assignment to the VA does not impact longer-term mortality risks, then quasi-experimental assignment of compliers implies that the mortality hazard rates for compliers assigned to the VA and those assigned to non-VA hospitals should be the same after $\bar{t}$. Visually, it appears that the potential hazard rates of the compliers are very similar in weeks $t \in \{1,\ldots,51\}$. Consistent with this impression, in online Appendix A.2, we show that we cannot reject that $h_{IV}(t;1) = h_{IV}(t;0)$ for all weeks $t \geq 1$, with a bootstrap-based p-value of 0.31. In contrast, in the next subsection, we show that longer-term mortality hazards are

[36] Specifically, testing equation (7) with respect to potential outcomes $y \in \mathcal{Y}$ may be more likely to detect violations of Condition 1 than the standard tests of monotonicity, focusing on patient characteristics, that we employ in online Appendix B. The intuition behind this is that the distributions of potential outcomes depend on both observed and unobserved patient characteristics, so violations in quasi-random assignment or monotonicity may be more likely to result in violations of equation (7). For example, if some ambulance companies take less sick veterans to the VA but sicker veterans to non-VA hospitals, we might expect to see implied negative mortality hazards for VA-assigned patients after a few weeks.

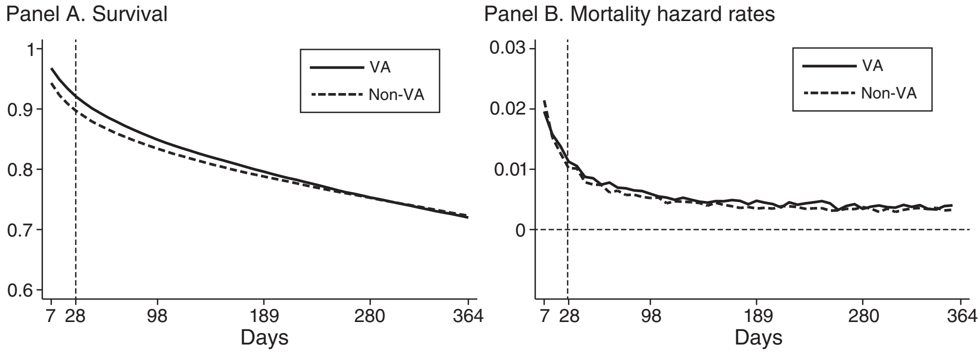

FIGURE 4. OBSERVED RISK-ADJUSTED OUTCOMES

Notes: This figure shows observed risk-adjusted outcomes for patients who arrive at a VA hospital and those who arrive at a non-VA hospital. Panel A shows survival outcomes as a function of days from the ambulance ride. "Days" indicate one-week intervals from the ambulance ride. Denote $S_i(t;d) \in \{0,1\}$ as an indicator for whether patient i survives up to time t after the ambulance ride, depending on whether the patient arrives at the VA ($d = 1$) or a non-VA hospital ($d = 0$). Observed survival is $S_i(t) = D_i S_i(t;1) + (1 - D_i) S_i(t;0)$. We estimate VA survival, or $E[S_i(t) \mid D_i = 1]$, by an OLS regression with a dependent variable of $S_i(t) D_i$ and the same design matrix implied by equation (2); we estimate non-VA survival, or $E[S_i(t) \mid D_i = 0]$, by a similar OLS regression with a dependent variable of $S_i(t)(D_i - 1)$. All regressions use a sample of ambulance rides with no prior ride in the last year and the same baseline controls as described in Figure 1. Panel B presents implied weekly mortality hazard rates, as given by equation (6).

persistently *higher* for the group of veterans who select the VA (i.e., always-takers and compliers for whom $D_i = 1$) compared to those who select a non-VA hospital (i.e., never-takers and compliers for whom $D_i = 0$).

D. *Selection and Differential Mortality Risks*

Finally, we take a closer look at death rates during the year after the ambulance ride to better understand the differences between our main OLS and IV estimates of the VA advantage. As shown in panel A of Figure 4, we find that, remarkably, OLS survival curves cross, at about nine to ten months after the ambulance ride. This crossing reflects a reversal in the sign of the OLS-estimated VA treatment effect: While patients arriving at the VA experience an immediate survival benefit, the survival benefit eventually reverses. Patients arriving at the VA are *more* likely to die within a year.

Consistent with this observed survival pattern, panel B of Figure 4 reveals a crossover in the observed mortality hazard rates for patients taken to VA versus non-VA hospitals—i.e., $h_{OLS}(t;1)$ and $h_{OLS}(t;0)$. In the first week after the ambulance ride, the death rate is lower for patients at the VA, though the gap between the VA and non-VA hazards is smaller than the corresponding potential-outcomes gap for compliers shown in Figure 3. After that, the hazard rate is consistently higher for patients at the VA than for those at a non-VA hospital (i.e., $h_{OLS}(t;1) > h_{OLS}(t;0)$ for $t \geq 1$).

This gap suggests differences in the long-term mortality risk between always-takers and never-takers, who are not quasi-experimentally assigned by ambulances to the VA versus non-VA hospitals. Thus, contrasting these results with those in Figure 3

Table 4—Effect of VA Hospitals on Other Outcomes

	Dependent variable				
	Admission (1)	Hospital days (2)	ED revisits (3)	Outpatient visits (4)	Spending (5)
Panel A. OLS					
VA hospital	−0.004	0.514	−0.036	0.200	840
	(0.003)	(0.045)	(0.007)	(0.017)	(87)
Outcome mean	0.589	4.380	0.318	1.443	12,265
Observations	401,319	401,319	401,319	401,319	401,319
Panel B. IV					
IV estimate	−0.090	−0.468	0.029	0.379	−2,598
	(0.032)	(0.434)	(0.044)	(0.174)	(820)
Outcome mean	0.589	4.380	0.318	1.443	12,265
Observations	401,319	401,319	401,319	401,319	401,319

Notes: This table shows OLS and IV estimates of the effect of VA hospitals on various outcomes. Hospital days count the number of inpatient days immediately following the ED visit; if the patient is not admitted, this equals zero for that visit. Outpatient visits count the number of VA and non-VA outpatient visits within one month of the ride. ED revisits count subsequent ED visits up to 14 days following the ride. Spending is defined as total spending over the 28 days following the ambulance ride. Panel A gives OLS estimates, $\hat{\beta}_{OLS}$, for β in equation (2). Panel B gives IV estimates, $\hat{\beta}_{IV}$, as well as the first-stage coefficient, $\hat{\pi}_1$ in equation (3), with respect to the leave-out probability of the assigned ambulance company to transport patients to the VA. All regressions use baseline controls, which are described in online Appendix Table A.2 and include patient zip code dummies, ALS/BLS dummies, source of the ambulance ride, time categories, and patient prior utilization. The estimation sample is described in online Appendix Table A.1.

for compliers, who are quasi-experimentally assigned, suggests an important selection pattern working against the VA advantage: Veterans with higher underlying long-term mortality risks select into VA hospitals. While the short-term VA advantage initially offsets this imbalance, higher mortality hazards among veterans at the VA reemerge soon after the first week. These differences in baseline mortality hazards accumulate over time to generate large differences in long-term survival.

To identify differences in the baseline mortality risk between VA-assigned compliers and always-takers, we compare $h_{IV}(t;1)$ and $h_{OLS}(t;1)$; to identify differences between non-VA-assigned compliers and never-takers, we compare $h_{IV}(t;0)$ and $h_{OLS}(t;0)$. In online Appendix A.2, we show that we cannot reject the null hypothesis that $h_{IV}(t;1) = h_{OLS}(t;1)$ for $t \geq 1$. However, we can strongly reject the null hypothesis that $h_{IV}(t;0) = h_{OLS}(t;0)$ for $t \geq 1$. The average value of $h_{IV}(t;0)$, for $t \geq 1$, is significantly larger than the corresponding average value of $h_{OLS}(t;0)$, for $t \geq 1$, which implies that never-takers are healthier than compliers. This survival analysis shows, with substantially more precision than that afforded by the benchmark analysis in Section IIC, that the VA advantage is larger than the (precisely estimated) OLS effect would imply. This strongly suggests that veterans who use the VA are sicker than veterans who do not and that our precise OLS estimate of the VA advantage in Section II is almost certainly a lower bound of the LATE.

IV. Mechanisms

This section probes further into the mechanisms behind the large VA mortality advantage. We divide our analyses into three parts. First, we explore heterogeneity in treatment effects along dimensions of patient and hospital characteristics. Second, we ask whether the VA produces superior health outcomes by spending more; spending less would imply mechanisms that improve productivity. Relatedly, we investigate differences in services reported by VA versus non-VA hospitals. Third, we indirectly assess the mechanisms of continuity of care, health IT, and integration of care among veterans who only use non-VA care.

A. *Heterogeneous Treatment Effects*

We explore several dimensions of potentially heterogeneous treatment effects. Overall, our evidence suggests a larger VA advantage among medically vulnerable veterans and those more likely to use the VA. However, the VA advantage holds broadly across all types of veterans, and it is not explained by characteristics of either VA or non-VA hospitals.

Complier Characteristics.—One explanation for $\hat{\beta}_{IV} > \hat{\beta}_{OLS}$ in Table 3 is that the complier population is very different from the overall population. So we begin by examining observable characteristics of compliers in our quasi-experiment relative to the overall sample.[37] Table 5 shows results for various characteristics. Compliers are more likely to be Black, to have lower income, and to have had a prior VA ED visit. Compliers have slightly fewer recorded Elixhauser comorbidities and are less likely to receive advanced life support. In online Appendix Table A.7, we show similar patterns when comparing always-takers and never-takers, following an approach from Dahl, Kostol, and Mogstad (2014) that we describe in online Appendix A.3. Consistent with our analysis in Section IIID, we find that VA always-takers have higher predicted mortality, based on observable characteristics, than either compliers or never-takers.[38]

Researchers and policymakers have noted a higher incidence of mental health and substance abuse issues among veterans (Tanielian and Jaycox 2008). Recognizing this need, in 1999 Congress allocated $152 million for increasing mental health care programming; in the following two decades, VA stations expanded mental health services and hired thousands of mental health providers (Veterans Millennium Health Care and Benefits Act 1999; US Government Accountability Office 2015). This capacity to treat mental health disorders contrasts with the non-VA health care sector, where mental health services have long been underfunded and underprovided (Huskamp and Iglehart 2016). We find higher rates of mental illness and substance abuse among compliers than in the overall population (Table 5); we similarly find higher rates of these conditions among always-takers than among never-takers (online Appendix Table A.7).

[37] Specifically, we employ the same approach from Abadie (2002) that we introduced in Section IIIA. Under IV validity in Condition 1, we can estimate $E[X_i \mid i \in \mathcal{C}]$ for some characteristic X_i by two-stage least squares, involving the first-stage equation (3) and a reduced-form equation replacing the outcome variable in equation (4) with $X_i D_i$.

[38] For Table 5 and online Appendix Table A.7, we predict mortality with baseline and holdout characteristics described in Section IIB. We do so using only rides going to non-VA hospitals, to separate the VA advantage from coefficients used in the prediction. An alternative approach is to estimate coefficients on predictors while controlling for D_i, which yields nearly identical results. To understand the rationale for this, consider a predictor that also predicts VA usage. The predictor may negatively predict mortality in the full sample when not conditioning on D_i, even if it positively predicts mortality conditional on D_i.

Table 5—Complier Characteristics

	Overall	Compliers	Ratio
Male	0.963	0.952	0.99
		(0.006)	[0.98–1.00]
Age	76.0	74.9	0.99
		(0.433)	[0.97–1.00]
Black	0.194	0.257	1.33
		(0.028)	[1.05–1.61]
Income	$20,905	$16,972	0.81
		($611)	[0.75–0.87]
Rural zip code	0.051	0.091	1.78
		(0.025)	[0.82–2.75]
Residential source	0.705	0.647	0.92
		(0.033)	[0.83–1.01]
Comorbidity count	6.143	5.447	0.89
		(0.113)	[0.85–0.92]
Mental illness	0.427	0.444	1.04
		(0.015)	[0.97–1.11]
Substance abuse	0.144	0.163	1.13
		(0.011)	[0.97–1.28]
Prior VA ED visit only	0.294	0.412	1.40
		(0.012)	[1.32–1.48]
Prior non-VA ED visit only	0.247	0.065	0.26
		(0.005)	[0.22–0.30]
Prior VA and non-VA ED visit	0.235	0.288	1.23
		(0.012)	[1.12–1.33]
Ambulance rides in prior year	2.156	2.178	1.01
		(0.084)	[0.93–1.09]
Advanced life support	0.684	0.600	0.88
		(0.024)	[0.81–0.95]
Predicted VA user	0.847	0.939	1.11
		(0.004)	[1.10–1.12]
Predicted mortality	0.097	0.095	0.98
		(0.004)	[0.90–1.06]

Notes: This table presents average complier characteristics and the ratio between this average and the average among all veterans in the sample. Average complier characteristics and standard errors are calculated by performing two-stage least squares using the first-stage equation (3) and a reduced-form equation replacing the outcome variable in equation (4) with $X_i D_i$, where X_i is the characteristic corresponding to ride i. Regressions use baseline controls described in online Appendix Table A.2; the regression sample is the baseline sample described in online Appendix Table A.1. Standard errors for each average are presented in parentheses. The corresponding 95 percent confidence intervals for each ratio are presented in brackets.

Selection Model.—Next, we explore whether compliers differ by their treatment effects, using the standard framework of an endogenous selection model. Following the "marginal treatment effects" (MTE) literature (see, e.g., Heckman and Vytlacil 2007 for a review), we exploit our multivalued ambulance instrument to characterize the relationship between treatment effects and the veteran's revealed propensity to go to the VA.

Specifically, we allow flexibility in the returns to VA care for compliers induced into VA care by ambulances along the continuum of VA shares. Compliers induced into VA care by ambulances with low VA shares reveal a higher propensity to use the VA than those who require ambulances with high VA shares to go to the VA. This approach allows compliers to differ—in both characteristics and potential outcomes—depending on the instrument (i.e., the quasi-experimentally assigned ambulance's VA share) through which they are induced into VA care. We provide further details of our approach in online Appendix A.4.

We find evidence of moderate "selection on gains," in which veterans with larger mortality reductions from going to the VA are more likely to go to the VA. In online Appendix Figure A.8, we show the MTE function ranging from veterans who are most likely to use the VA to those who are least likely to use the VA. Veterans induced to go to the VA by lower-propensity ambulances have higher returns to VA care than those induced by high-propensity ambulances. In online Appendix Table A.8, we find a substantial ATE, only marginally smaller than the LATE, across various specifications.

Heterogeneity by Patient and Hospital Characteristics.—Finally, we assess heterogeneity in the VA advantage by observable hospital and patient characteristics. We consider characteristics in three categories: (i) patient characteristics; (ii) characteristics of non-VA hospitals serving a given zip code, weighting the hospitals by volume of rides from the zip code; and (iii) characteristics of the VA hospital serving a given zip code. Online Appendix A.5 provides details on hospital characteristics; online Appendix A.6 describes our estimation approach.

Online Appendix Table A.9 shows results for patient characteristics. The VA advantage is substantially larger for medically vulnerable veterans. Veterans with higher predicted mortality, those transported by ambulances offering advanced life support, and those with more ambulance rides in the prior year have larger treatment effects. The VA survival benefit appears greater for veterans suffering from mental illness or substance abuse and for those with greater prior utilization in the VA. However, none of the differences in the VA survival benefit across patient characteristics imply a group harmed by the VA. Notably, the VA survival benefit is not limited to select medical conditions that stereotypical users of the VA might have; even patients who are less likely to use VA care experience a similar VA survival benefit.

Table 6 shows differences in hospital characteristics between VA and non-VA hospitals. For example, VA hospitals have fewer ED visits and admissions per bed and are more likely to be teaching hospitals.[39] However, we find only modest treatment heterogeneity with respect to any of these hospital characteristics at the zip code level (online Appendix Tables A.10 to A.13).[40] Heterogeneity along any of the VA or non-VA hospital characteristics across zip codes is less than 20 percent of the main VA advantage, suggesting that the VA advantage remains across the spectrum of VA and non-VA alternatives. Zip codes from which non-VA rides predominantly go to a single (non-VA) hospital have a smaller VA advantage. As for VA hospital characteristics, the VA advantage is greater for larger VA hospitals. In

[39] VA hospitals appear to have more long-term care admissions, which explains a higher average length of stay (i.e., fewer admissions for a slightly larger average daily census). As shown in Table 4, the difference in length of stay is not borne out in our sample; the IV estimate of the effect on length of stay suggests that the VA reduces length of stay.

[40] To assign VA hospital characteristics to a given zip code, we use the characteristics of the VA hospital that receives the highest share of ambulance rides originating from that zip code. For non-VA hospital characteristics, we average characteristics across non-VA hospitals, weighted by the number of rides going to each non-VA hospital from a given zip code. Finally, we consider the difference between the two by subtracting the (average) characteristic for the non-VA hospitals from the corresponding characteristic of the assigned VA hospital.

Table 6—Means of Hospital Characteristics

	Hospital sample			
	VA		Non-VA	
	Baseline sample	AHA average	Baseline sample	AHA average
Volume, size, and capabilities				
ED visits	17,287	20,664	41,600	38,416
Admissions	6,234	7,947	13,676	12,445
Average daily census	224	315	200	182
Total staffed beds	317	446	283	258
Teaching hospital	0.59	0.58	0.27	0.22
Urban location	0.84	0.89	0.97	0.89
Trauma center	0.10	0.13	0.64	0.61
Advanced cardiac care	0.71	0.73	0.70	0.58
Staffing				
ED staff per 1,000 ED visits	0.84	0.79	0.55	0.00
Nurses per 1,000 patient-days	6.07	5.77	5.45	0.01
Physicians per 1,000 patient-days	4.93	4.89	7.75	0.01
Hospitalists per 1,000 patient-days	0.15	0.17	0.29	0.00
Intensivists per 1,000 patient-days	0.08	0.06	0.16	0.00
Spending and relative outcomes				
Relative spending	1.12	1.10	1.01	1.00
Mortality rate	7.69	7.58	12.27	12.23
Readmissions rate	12.38	12.23	18.08	18.14
Payment, organization, and health IT				
Capitated lives covered			8,087	11,399
Network participant			0.51	0.46
Hospital system			0.75	0.62
HMO			0.19	0.20
PPO			0.21	0.19
ACO			0.04	0.09
Adopted health IT by 2011			0.25	0.20
Adopted health IT by 2014			0.93	0.87
Hospital count	98	122	1,385	8,845

Notes: This table presents average characteristics of VA and non-VA hospitals in the sample of hospitals included in the American Hospital Association (AHA) Annual Survey and in the subset of these hospitals that are in our baseline sample. Non-VA hospital characteristics are further presented for the baseline sample and for the national average. The national average weights hospital characteristics by their number of yearly admissions as recorded in the American Hospital Association Annual Survey. The average in the baseline sample weights hospital characteristics by rides in that sample, described in online Appendix Table A.1. Relative spending measures are not comparable between VA and non-VA hospitals. Hospital characteristics are described in further detail in online Appendix A.5.

online Appendix A.7, we describe complementary results from an empirical Bayes approach to heterogeneity in station-specific OLS estimates of the VA advantage; in that approach, we fail to demonstrate meaningful heterogeneity in the VA advantage across VA stations.

B. *Effect on Spending and Utilization*

In light of the important literature on the returns to spending in health care (e.g., Garber and Skinner 2008), we examine the causal effect of VA versus non-VA care on spending. The motivation behind this analysis is similar to that in Doyle et al.

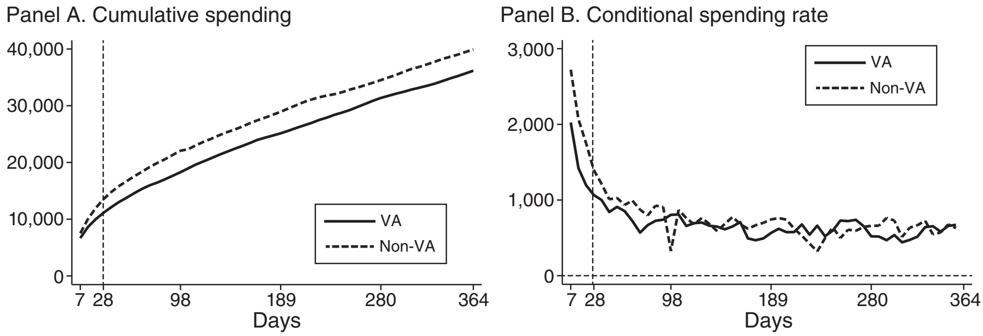

FIGURE 5. COMPLIER SPENDING

Notes: This figure shows potential spending outcomes for ambulance compliers who arrive at a VA hospital and those who arrive at a non-VA hospital. Denote $\text{Spending}_i(t;d)$ as the potential cumulative spending function for patient i up to time t after the ambulance ride, depending on whether the patient arrives at the VA $(d = 1)$ or a non-VA hospital $(d = 0)$. If a veteran i dies at $\underline{t}$, $\text{Spending}_i(t;d)$ will be constant for all $t \geq \underline{t}$. Panel A shows cumulative spending per patient as a function of days from the ambulance ride. We estimate cumulative spending for compliers who arrive at a VA hospital, i.e., $E[\text{Spending}_i(t;1) \mid i \in \mathcal{C}]$, by an IV regression with a dependent variable of $\text{Spending}_i(t) \times D_i$, the endogenous VA treatment D_i, and the same first-stage and reduced-form design matrix implied by equations (3) and (4). We estimate complier non-VA cumulative spending, i.e., $E[\text{Spending}_i(t;0) \mid i \in \mathcal{C}]$, by a similar IV regression with a dependent variable of $\text{Spending}_i(t)(D_i - 1)$. All regressions use a sample of ambulance rides with no prior ride in the last year and the same baseline controls as described in Figure 1. Panel B presents implied weekly spending rates for compliers, conditional on survival. See online Appendix Figure A.10 for spending results with prices fixed.

(2015), who sought to understand whether higher-spending hospitals achieve better health outcomes. We also move beyond aggregate spending to examine the nature of services reportedly delivered in VA and non-VA care.

Actual Spending.—We calculate our baseline measure of spending from the perspective of actual spending by taxpayers and veterans, relying on internal VA cost data and Medicare payment data from claims. Internal VA cost accounting apportions costs by VA utilization data and scales the cost of each encounter so that total spending matches actual budgeted spending within each VA station.[41] On the Medicare side, we include total payments made to providers, including those from the veteran (i.e., coinsurance and deductible), the government, and any other insurer. This measure accords with the government's view of resources—paid for by taxpayers and consumers—to procure or provide health services. Differences in resources could be driven by differences between the VA and Medicare in the quantity of services provided and the "prices" for delivering a given service.

Using the same instrumental variables approach as in our benchmark analysis, we study the VA effect on spending over time since the ambulance ride. Specifically, we combine VA and Medicare spending in various weekly intervals since the ambulance ride. Table 4 further shows that the VA reduces 28-day combined spending

[41] The apportioning uses inputs such as Relative Value Units (RVUs) associated with CPT codes, Diagnosis-Related Group (DRG) weights, patient characteristics, and admission lengths of stay. This methodology is detailed in Wagner, Chen, and Barnett (2003) and in Phibbs et al. (2019).

by $2,598, or 21 percent of the mean 28-day spending. The reduction in spending reflects a lower probability of inpatient admission and fewer hospital days associated with VA care. The corresponding OLS regressions of these outcomes are less favorable for the VA: VA arrival is associated with more hospital days and higher spending. Consistent with our earlier findings, this suggests that VA always-takers require more care than never-takers. Interestingly, the VA does not uniformly reduce utilization—the VA *increases* outpatient visits in the following 28 days. Figure 5 shows potential cumulative spending curves during the first year and implied weekly spending rates conditional on survival, for compliers transported to a VA hospital and those transported to a non-VA hospital. Differences in cumulative spending accrue until about three months after the ambulance ride; thereafter, the differences remain stable.

Fixed-Price Spending and Reported Utilization.—As an alternative measure of spending, we take reported utilization in the VA and Medicare data and apply the same prices to each instance of utilization, based on its identifying code, regardless of whether it was covered by the VA budget or reimbursed by Medicare.[42] By construction, this measure focuses on differences in the quantities of each reported service between the VA and Medicare. According to this measure, the VA reduces 28-day combined spending by much more: $5,267, about double the reduction in our baseline measure of actual spending.[43]

Importantly, both VA and non-VA hospitals use identical coding systems for recording utilization and are held to the same standard of accurate coding based on clinical documentation. However, financial incentives in the two settings differ starkly. Outside of the VA, financing is predominantly fee-for-service, tightly connected with units of utilization, and based on notions of cost. Inside the VA, financing is based on the population of veteran enrollees and much less connected with reported utilization (Wasserman et al. 2005).

We uncover stark differences in reported 28-day utilization between the VA and non-VA hospitals.[44] These differences likely reflect a combination of differences in reporting (Dafny 2005; Fang and Gong 2017) and differences in actual utilization. In Figure 6, we examine the most common 25 CPT codes in our sample and show wide variation in their utilization at the VA relative to outside of the VA. Some outpatient and rehabilitation services (e.g., CPT codes 99212 and 99110) are much more likely to be performed in the VA than outside of it; remarkably, telephone calls (CPT code 98966) are *only* reported at the VA. In Figure 7, we show that services more likely to be performed in non-VA hospitals have much higher reimbursement

[42] We closely follow methodology laid out by Gottlieb et al. (2010) and Finkelstein, Gentzkow, and Williams (2016). Specifically, we impute spending for physician services based on Relative Value Units for service procedures with CPT codes, for other outpatient procedures based on average reimbursements for (non-CPT) HCPCS codes, and for inpatient stays based on Diagnosis-Related Group weights obtained from the National Bureau of Economic Research (1984-2023). Data on CPT and HCPCS payment-related modifiers and Medicare inpatient payment base rates were manually compiled from public sources, as detailed in our replication package (Chan, Card, and Taylor 2023). We scale prices by a constant so that imputed total Medicare spending equals actual total Medicare spending.

[43] See online Appendix Figure A.10 for spending potential outcomes when considering spending with prices held fixed between VA and Medicare utilization. This figure follows Figure 5 in format.

[44] These results remain qualitatively similar when we examine reported seven-day utilization instead.

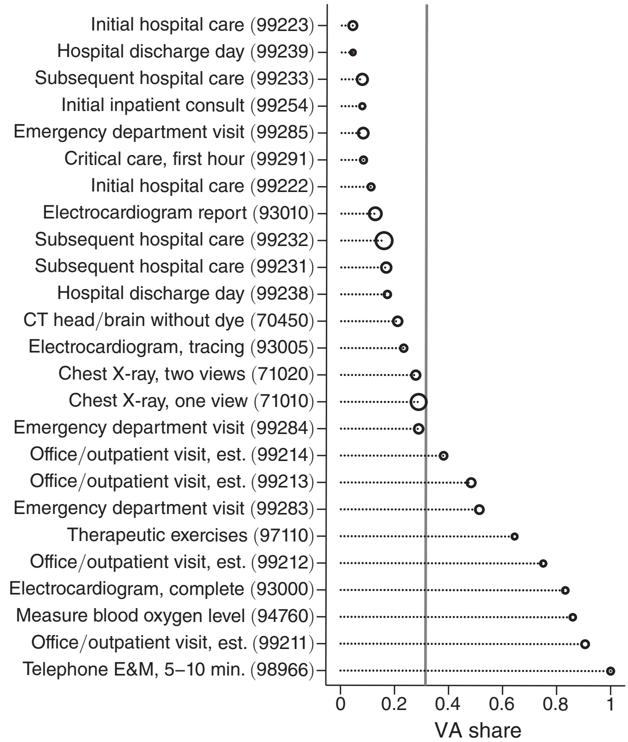

FIGURE 6. VA SHARES WITHIN TOP REPORTED PROCEDURES

Notes: This figure shows the VA share of reported utilization of each of the top 25 procedures, defined by Current Procedural Terminology (CPT) codes, on the Medicare Physician Fee Schedule (MPFS). We include utilization for any patient in our baseline sample in the 28 days following their ambulance ride. The area of each circle indicates the relative utilization volume of each CPT code. For scale, the largest circle represents a service utilization of 1.627 times per ambulance ride; the smallest circle represents a service utilization of 0.162 times per ambulance ride. The top 25 procedures represent 57.2 percent of the total utilization of all CPT codes on the MPFS. The gray vertical line indicates the overall VA share of utilization of any CPT code on the MPFS.

under Medicare (and other fee-for-service arrangements); in contrast, services that are more common in the VA receive very little reimbursement. In online Appendix Figure A.12, we show that, for evaluation and management (E/M) services that are more highly reimbursed when patients are reported as being more complex, the odds of reporting high versus low complexity are five times higher in private hospitals versus the VA. Online Appendix A.8 describes these analyses in further detail.

Implications.—The result that the VA saves lives while reducing spending is significant for two reasons. First, the result speaks directly to the policy question of whether the VA should privatize its care in a Medicare-type arrangement. We show that, at least for the patients in our design, this privatization arrangement would be dominated by the status quo, as it would lead to both higher spending and worse health outcomes. Second, this joint finding suggests that the general mechanism behind the VA survival benefit is not higher spending but higher productivity.

Our evidence points to productive inefficiency, rather than "flat of the curve" spending, underlying the relatively low returns to US health care. We also uncover

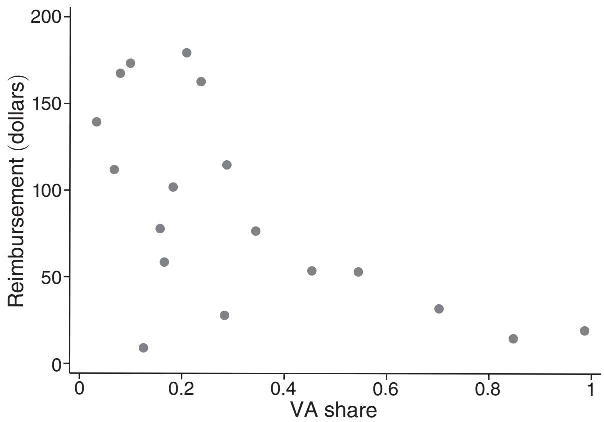

Figure 7. Reimbursement and VA Shares of Procedures

Notes: This figure shows a binned scatterplot between reimbursement in dollars (on the y-axis) and the VA share of utilization corresponding to each of the top 200 Current Procedural Terminology (CPT) codes on the Medicare Physician Fee Schedule (MPFS). Reimbursement for a CPT code is calculated for each year it is on the MPFS by multiplying its year-specific relative value units (RVUs) with the year-specific conversion factor. We then take the median reimbursement across years that the CPT code is on the MPFS. We include utilization for any patient in our baseline sample in the 28 days following their ambulance ride. The top 25 procedures represent 79.3 percent of the utilization of any CPT code on the MPFS. We weight each CPT observation by its utilization in forming the binned scatterplot.

striking differences in reported utilization between the VA and non-VA sectors. Public versus private provision of care implies fundamentally different sets of financial incentives, which may plausibly drive stark differences in both reported and actual utilization. The private sector neglects many services with low reimbursement, but these services may nonetheless improve coordination of care and health outcomes. The private sector reliance on fee-for-service billing may also imply differences in the share of time physicians spend on documentation and differences in the preferences and skills of physicians who choose to work in the respective environments.

These results complement a growing literature on productivity differences across personnel (Chan, Gentzkow, and Yu 2022; Silver 2020) and hospitals (Chandra and Staiger 2007, 2020) by showing an important productivity difference between health care *systems*. They also add to our understanding of the wide variation in spending well documented by Medicare claims within the private sector (e.g., Fisher et al. 2003; Finkelstein, Gentzkow, and Williams 2016).[45] Finally, these results shed light on how the privatization or "corporatization" of health care in the US fee-for-service context may increase profitability while worsening patient health outcomes (Eliason et al. 2020) and restricting access for disadvantaged patients (Duggan et al. 2023).

[45] Fisher et al. (2003) study a national cohort of Medicare beneficiaries and show that patients who live in hospital referral regions in the highest-quintile of end-of-life spending have more than double the inpatient visits, tests, and minor procedures than patients who live in the lowest-quintile HRRs.

C. *Health IT and Integrated Care*

Our final analysis investigates the role of health IT and integrated care in generating the VA survival advantage. There is much in the qualitative literature to support this mechanism. Fragmentation and poor coordination in the US health sector have long been highlighted as potential sources of inefficiency. In the VA, a qualitative literature attributes its "transformation" into a high-quality health system, achieving superior performance in a wide range of process measures, to its adoption of health IT and integrated care in the mid-1990s (e.g., Jha et al. 2003).[46]

While empirical research has focused on one mechanism or another, key complementarities likely exist between a patient's continuity of care and a health system's adoption of health IT and integrated care. Patients with isolated problems or those with no prior utilization in a health care system will benefit little from health IT and integrated care. Similarly, patients who are well informed or whose providers have time to make decisions (e.g., hip replacements, chemotherapy initiation) may overcome informational barriers imposed by the lack of health IT or integrated care. In contrast, easy access to information is likely crucial in our setting of emergency conditions.

In this subsection, we can provide only indirect evidence of these mechanisms. We cannot assess the VA's effect prior to its implementation of health IT and its reorganization into more integrated care, since these events in the mid-1990s predate the availability of data for analysis.[47] Similarly, we cannot examine the VA's effect on mortality among veterans who would not benefit from continuity of care, since ambulances rarely transport veterans without prior VA utilization to the VA, as shown in online Appendix Figure A.13. Thus, we will examine these mechanisms in a separate sample of veterans who only use providers outside of the VA. As online Appendix Figure A.13 also shows, veterans may utilize more than one non-VA hospital system; ambulances may or may not transport these veterans to the hospital system where they usually receive care.

We detail our analysis in online Appendix A.9. In brief, we construct a separate sample of veterans who have only used non-VA care in the prior year and live in a zip code with more than one nearby non-VA hospital. Ambulances will almost certainly transport these veterans to a non-VA hospital. We use an ambulance instrument similar to the one we use in our benchmark analysis to assess mortality outcomes depending on whether these veterans are quasi-randomly assigned to their most commonly visited (non-VA) hospital in the prior year (i.e., their "modal" hospital). This modal hospital effect on mortality arguably captures at least some of the potential effect of continuity of care in the private sector. We then examine complementarities with health IT and integrated care by further exploiting two changes induced by incentives in federal laws and payment policies during our study period. First, the HITECH Act of 2009 dramatically increased the share of hospitals using health IT (Blumenthal 2010).[48] Second, in

[46] For recent qualitative research that illuminates this mechanism between VA and non-VA care, see, e.g., Nevedal et al. (2019) and Rinne et al. (2019).

[47] Indeed, the VA's adoption of a standardized health IT platform (VistA) in the mid-1990s paved the way for research on health services within the VA system, including this study.

[48] In 2009, 1.5 percent of US nonfederal hospitals had an electronic health record (EHR) system in all clinical units, and an additional 7.6 percent had an EHR system in at least 1 clinical unit (Jha et al. 2009). By 2014, 97 percent of such hospitals possessed an EHR technology meeting requirements of the Department of Health and Human Services, and 76 percent of hospitals had implemented the EHR system in at least 1 clinical unit (Charles, Gabriel,

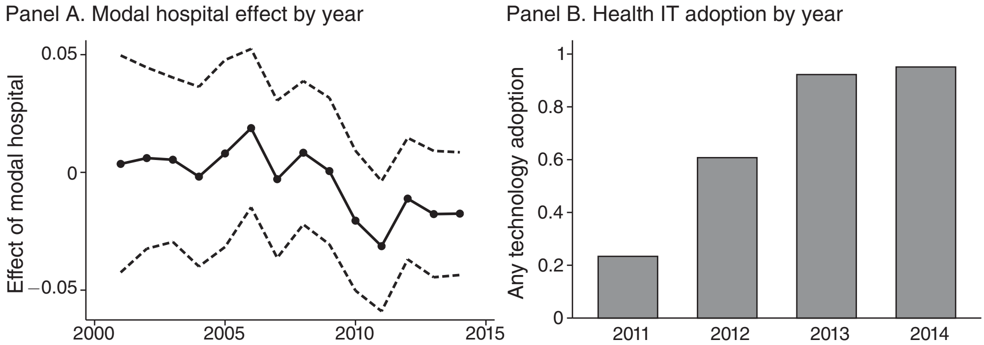

Figure 8. Modal Hospital Effect and Health IT Adoption

Notes: Panel A of this figure shows the IV estimate of the modal non-VA hospital effect on 28-day mortality by calendar year. The first-stage and reduced-form equations are given in equations (A.30) and (A.31) in the online Appendix. The overall sample is described in online Appendix Table A.14 and contains ambulance rides for veterans with no prior VA use in the previous year. Results for the overall IV estimates are shown in online Appendix Figure A.14. Estimates are shown in connected dots, while 95 percent confidence intervals are shown in dashed lines. Panel B of the figure shows the percent of rides going to hospitals after health IT adoption in our analytic sample. Health IT adoption is defined from a dataset from the Office of the National Coordinator of Health Information Technology (ONC). This dataset merges hospital attestation data from the Medicare EHR Incentive Program with certified EHR product information from ONC's Certified Health IT Product List (CHPL), and we code the use of any certified product as health IT adoption.

2011, Medicare began to incentivize care integration via alternative payment arrangements to "Accountable Care Organizations" (Greaney 2011).

We find that, overall, the survival advantage of being transported to a veteran's modal hospital is small. Importantly, however, we show in Figure 8 that the effect only begins to appear around the time of the HITECH Act of 2009, when a large share of non-VA hospitals adopted health IT. We also examine how the modal hospital effect relates to hospital-specific dates of health IT or ACO adoption. Our results provide suggestive evidence that the growth in the modal hospital effect is associated with health IT adoption, even when holding hospitals fixed. The relationship with ACO adoption appears similar but is imprecise (see online Appendix Table A.16).[49] Nonetheless, even after 2009, the modal hospital effect among private hospitals is at most half the size of the VA advantage. This suggests important and complementary mechanisms outside of health IT and integrated care (e.g., financial incentives), as described in Section IVA, or differences in the VA's implementation of health IT and integrated care.[50]

and Searcy 2015). However, interoperability (i.e., the ability to share electronic medical records) across private hospitals has remained low (Holmgren, Patel, and Adler-Milstein 2017). To this day, multiple EHR platforms exist in the private sector, and they do not communicate with each other.

[49] In contrast to dramatic rates of health IT adoption, we find that only 11 percent of non-VA hospitals participated in ACOs by the last year of our sample, consistent with other research (Colla et al. 2016).

[50] This result supports Starr's (2008) observation, "Using technology to fix problems of organization has been a perennial hope in medical care and, as in many other areas, a perennial disappointment when the efforts amount only to automating the status quo. Combining changes in organization, payment, and technology has more promise. Although the ultimate impact of these combined efforts is unclear, they have already gone from pilot projects to significant changes in national policy" (p. 487).

V. Conclusion

The structure of health care delivery to US veterans provides a distinctive research opportunity, allowing us to study fundamentally different systems of health care that coexist for a large patient population. Specifically, millions of older veterans (those at least age 65) are dually eligible for care in a public system operated by the Veterans Health Administration or in private sector hospitals financed by Medicare. The ambulance setting provides plausible quasi-experimental assignment of veterans to these health care systems. Our work has current policy relevance, as the Department of Veterans Affairs is now considering whether to bolster its existing public delivery system or to replace it, either partially or entirely, with a system of financing private care.

We find a significant VA advantage: Our preferred instrumental variables estimate, based on veterans induced by their ambulance company to be treated at a VA or non-VA hospital, shows a 4.5 pp survival gain at 28 days (95 percent confidence interval 1.1 to 8.0 pp), implying about a 46 percent reduction in mortality relative to the overall average. These survival gains occur in the first week following the ambulance ride and appear to be long-lasting. We further use this survival-analysis framework to validate our IV quasi-experiment and to demonstrate differences in long-term mortality hazards between VA and non-VA users who are noncompliers; we show that, if anything, the VA treats patients in our sample who are sicker than the ones non-VA hospitals treat. Although we find some intuitive margins of heterogeneity in the VA advantage, the VA outperforms the non-VA alternatives in a wide variety of locations and for all types of patients we consider, not only for patients with stereotypical medical conditions.

Importantly, the VA also reduces total spending by 21 percent relative to non-VA providers, which points to higher productivity in the VA than in the private sector. We demonstrate striking differences in the procedures reportedly performed at the VA compared to private providers. These differences relate to the underlying arrangements by which public versus private providers are funded in the US (and other developed countries). We also find evidence consistent with complementary mechanisms of continuity of care, health IT, and organization. For example, veterans with prior VA care (and those who are more likely to use the VA) have larger survival gains from VA assignment. Among veterans who only use non-VA hospitals, the benefits of continuity of care are weaker and seem to materialize only when their (non-VA) hospitals adopt health IT and integrated care. Still, even when accounting for private efforts to adopt health IT and reorganize care, a sizable residual VA advantage remains.

Our results contribute more broadly to two streams of literature on the efficiency of production. First, we contribute to the descriptive analysis that compares the performance of the US health care system to systems in other developed countries (Blank, Burau, and Kuhlmann 2017). By almost all accounts, comparisons of US health outcomes and health care spending are unfavorable with those of other developed countries (Garber and Skinner 2008; Rice et al. 2013). Our analysis points to a potentially significant source of inefficiency in the US context: its version of private health care provision. This arrangement rewards costly but not necessarily efficient care. Although several developed countries that outperform the United States also feature private provision of care, the US system arguably has the most complex configuration of financing and delivery, with high levels of uninsurance, administrative costs, and fragmentation (Cebul et al. 2008). These well-known information and coordination gaps may be fatal, at least for veterans in emergencies.

Second, we provide empirical support, in the context of health care, for the general idea of production complementarities among three innovations in production: workplace reorganization, products and services, and information technology (Bresnahan, Brynjolfsson, and Hitt 2002). The VA adopted a comprehensive health IT system almost two decades before nearly all private hospitals in the United States. This reform was accompanied by a massive integration of care, reorganizing the delivery system and redefining services involved in patient care. For private hospitals, redefining health care products and services is limited by fee-for-service payment systems and the difficulty of measuring quality (Cutler 2010). Hospitals without a broad network of clinics and a clear mandate for a population's health may find it difficult to reorganize and redefine their services to optimize patient health. Our result that health IT in private hospitals may improve survival—but to a muted extent and only for patients the hospitals have previously treated—is consistent with these production complementarities. Complementarities in health care production may pose barriers for replicating the VA advantage but also argue for coordinated policies in the fragmented private landscape of US health care.

REFERENCES

Abadie, Alberto. 2002. "Bootstrap Tests for Distributional Treatment Effects in Instrumental Variable Models." *Journal of the American Statistical Association* 97 (457): 284–92.

Abaluck, Jason, Mauricio Caceres Bravo, Peter Hull, and Amanda Starc. 2021. "Mortality Effects and Choice Across Private Health Insurance Plans." *Quarterly Journal of Economics* 136 (3): 1557–1610.

Agha, Leila. 2014. "The Effects of Health Information Technology on Costs and Quality of Medical Care." *Journal of Health Economics* 34: 19–30.

Agha, Leila, Brigham Frandsen, and James B. Rebitzer. 2019. "Fragmented Division of Labor and Healthcare Costs: Evidence from Moves Across Regions." *Journal of Public Economics* 169: 144–59.

Agha, Zia, Richard P. Lofgren, Jerome V. VanRuiswyk, and Peter M. Layde. 2000. "Are Patients at Veterans Affairs Medical Centers Sicker? A Comparative Analysis of Health Status and Medical Resource Use." *Archives of Internal Medicine* 160 (21): 3252–57.

Alchian, Armen A. 1965. "Some Economics of Property Rights." *Il Politico* 30 (4): 816–29.

Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber. 2005. "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools." *Journal of Political Economy* 113 (1): 151–84.

Altonji, Joseph G., and Richard K. Mansfield. 2018. "Estimating Group Effects Using Averages of Observables to Control for Sorting on Unobservables: School and Neighborhood Effects." *American Economic Review* 108 (10): 2902–46.

Altonji, Joseph G., and Lewis M. Segal. 1996. "Small-Sample Bias in GMM Estimation of Covariance Structures." *Journal of Business & Economic Statistics* 14 (3): 353–66.

American Medical Association. 2015. *The Physicians' Guide: Medicare RBRVS.* Chicago, IL: American Medical Association.

Andreyeva, Elena, Atul Gupta, Catherine Ishitani, Malgorzata Sylwestrzak, and Benjamin Ukert. 2022. "The Corporatization of Hospital Care." SSRN Working Paper 4134007.

Angrist, Joshua D., Peter D. Hull, Parag A. Pathak, and Christopher R. Walters. 2017. "Leveraging Lotteries for School Value-Added: Testing and Estimation." *Quarterly Journal of Economics* 132 (2): 871–919.

Arnold, David, Will Dobbie, and Crystal S. Yang. 2018. "Racial Bias in Bail Decisions." *Quarterly Journal of Economics* 133 (4): 1885–1932.

Balke, Alexander, and Judea Pearl. 1997. "Bounds on Treatment Effects from Studies with Imperfect Compliance." *Journal of the American Statistical Association* 92 (439): 1171–76.

Bhuller, Manudeep, Gordon B. Dahl, Katrine V. Loken, and Magne Mogstad. 2020. "Incarceration, Recidivism, and Employment." *Journal of Political Economy* 128 (4): 1269–1324.

Blank, Robert, Viola Burau, and Ellen Kuhlmann. 2017. *Comparative Health Policy.* 5th ed. London: Red Globe Press.

Blom, David. 2016. *Considerations and Recommendations*. Washington, DC: Commission on Care.

Bloom, Nicholas, Carol Propper, Stephan Seiler, and John Van Reenen. 2015. "The Impact of Competition on Management Quality: Evidence from Public Hospitals." *Review of Economic Studies* 82 (2): 457–89.

Blumenthal, David. 2010. "Launching HITECH." *New England Journal of Medicine* 362 (5): 382–85.

Blundell, Richard, and Rosa L. Matzkin. 2014. "Control Functions in Nonseparable Simultaneous Equations Models." *Quantitative Economics* 5 (2): 271–95.

Bohmer, Richard M. J., Amy C. Edmondson, and Laura Feldman. 2013. "Intermountain Health Care." Harvard Business School Case 603–066.

Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt. 2002. "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence." *Quarterly Journal of Economics* 117 (1): 339–76.

Card, David, Carlos Dobkin, and Nicole Maestas. 2007. "Does Medicare Save Lives?" NBER Working Paper 13668.

Card, David, Carlos Dobkin, and Nicole Maestas. 2009. "Does Medicare Save Lives?" *Quarterly Journal of Economics* 124 (2): 597–636.

Card, David, Alessandra Fenizia, and David Silver. 2023. "The Health Impacts of Hospital Delivery Practices." *American Economic Journal: Economic Policy* 15 (2): 42–81.

Cebul, Randall D., James B. Rebitzer, Lowell J. Taylor, and Mark E. Votruba. 2008. "Organizational Fragmentation and Care Quality in the U.S. Healthcare System." *Journal of Economic Perspectives* 22 (4): 93–113.

Centers for Medicare and Medicaid Services. 2002. "Medicare Program; Fee Schedule for Payment of Ambulance Services and Revisions to the Physician Certification Requirements for Coverage of Nonemergency Ambulance Services." In *Federal Register* 67 (39): 9099–9135.

Chan, David C., David Card, and Lowell Taylor. 2023. "Replication data for: Is There a VA Advantage? Evidence from Dually Eligible Veterans." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.3886/E192055V1.

Chan, David, Mark Duggan, and Audrey Guo. 2021. *Soldiering On: Improving Policies to Benefit America's Veterans*. Stanford, CA: Stanford Institute for Economic Policy Research.

Chan, David C., Matthew Gentzkow, and Chuan Yu. 2022. "Selection with Variation in Diagnostic Skill: Evidence from Radiologists." *Quarterly Journal of Economics* 137 (2): 729–83.

Chandra, Amitabh, and Douglas O. Staiger. 2007. "Productivity Spillovers in Healthcare: Evidence from the Treatment of Heart Attacks." *Journal of Political Economy* 115 (1): 103–40.

Chandra, Amitabh, and Douglas O. Staiger. 2020. "Identifying Sources of Inefficiency in Healthcare." *Quarterly Journal of Economics* 135 (2): 785–843.

Chandra, Amitabh, Amy Finkelstein, Adam Sacarny, and Chad Syverson. 2016. "Healthcare Exceptionalism? Productivity and Allocation in the U.S. Healthcare Sector." *American Economic Review* 106 (8): 2110–44.

Charles, Dustin, Meghan Gabriel, and Talisha Searcy. 2015. *Adoption of Electronic Health Record Systems among U.S. Non-Federal Acute Care Hospitals: 2008–2014*. Washington, DC: Office of the National Coordinator for Health Information Technology.

Chiang, Arthur J., Guy David, and Michael Gene Housman. 2006. "The Determinants of Urban Emergency Medical Services Privatization." *Critical Planning* 13.

Colla, Carrie H., Valerie A. Lewis, Emily Tierney, and David B. Muhlestein. 2016. "Hospitals Participating in ACOs Tend to Be Large and Urban, Allowing Access to Capital And Data." *Health Affairs* 35 (3): 431–39.

Cooper, Zack, Stephen Gibbons, and Matthew Skellern. 2018. "Does Competition from Private Surgical Centres Improve Public Hospitals' Performance? Evidence from the English National Health Service." *Journal of Public Economics* 166: 63–80.

Cubanski, Juliette, Tricia Neuman, Anthony Damico, and Karen Amith. 2018. *Medicare Beneficiaries' Out-of-Pocket Health Care Spending as a Share of Income Now and Projections for the Future*. San Francisco, CA: Kaiser Family Foundation.

Cutler, David. 2010. "Where Are the Health Care Entrepreneurs?" *Issues in Science and Technology* 27 (1): 49–56.

Dafny, Leemore S. 2005. "How Do Hospitals Respond to Price Changes?" *American Economic Review* 95 (5): 1525–47.

Dahl, Gordon B., Andreas Ravndal Kostol, and Magne Mogstad. 2014. "Family Welfare Cultures." *Quarterly Journal of Economics* 129 (4): 1711–52.

Department of Veterans Affairs. 2020. *FY 2021 Budget Submission: Budget in Brief*. Washington, DC: Department of Veterans Affairs.

Deryugina, Tatyana, and David Molitor. 2020. "Does When You Die Depend on Where You Live? Evidence from Hurricane Katrina." *American Economic Review* 110 (11): 3602–33.

Doyle, Joseph J. 2011. "Returns to Local-Area Health Care Spending: Evidence from Health Shocks to Patients Far from Home." *American Economic Journal: Applied Economics* 3 (3): 221–43.

Doyle, Joseph J., John A. Graves, Jonathan Gruber, and Samuel Kleiner. 2015. "Measuring Returns to Hospital Care: Evidence from Ambulance Referral Patterns." *Journal of Political Economy* 123 (1): 170–214.

Duggan, Mark, Atul Gupta, Emilie Jackson, and Zachary S. Templeton. 2023. "The Impact of Privatization: Evidence from the Hospital Sector." NBER Working Paper 30824.

Eliason, Paul J., Benjamin Heebsh, Ryan C. McDevitt, and James W. Roberts. 2020. "How Acquisitions Affect Firm Behavior and Performance: Evidence from the Dialysis Industry." *Quarterly Journal of Economics* 135 (1): 221–67.

Elixhauser, Anne, Claudia Steiner, D. Robert Harris, and Rosanna M. Coffey. 1998. "Comorbidity Measures for Use with Administrative Data." *Medical Care* 36 (1): 8–27.

Fang, Hanming, and Qing Gong. 2017. "Detecting Potential Overbilling in Medicare Reimbursement via Hours Worked." *American Economic Review* 107 (2): 562–91.

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams. 2016. "Sources of Geographic Variation in Health Care: Evidence from Patient Migration." *Quarterly Journal of Economics* 131 (4): 1681–1726.

Finkelstein, Amy, Matthew Gentzkow, and Heidi Williams. 2021. "Place-Based Drivers of Mortality: Evidence from Migration." *American Economic Review* 111 (8): 2697–2735.

Fisher, Elliott S., David E. Wennberg, Therese A. Stukel, Daniel J. Gottlieb, F. L. Lucas, and Etoile L. Pinder. 2003. "The Implications of Regional Variations in Medicare Spending. Part 1: The Content, Quality, and Accessibility of Care." *Annals of Internal Medicine* 138 (4): 273–87.

Frakes, Michael D., Jonathan Gruber, and Timothy Justicz. 2020. "Public and Private Options in Practice: The Military Health System." NBER Working Paper 28256.

Garber, Alan M., and Jonathan Skinner. 2008. "Is American Health Care Uniquely Inefficient?" *Journal of Economic Perspectives* 22 (4): 27–50.

Gaynor, Martin, Rodrigo Moreno-Serra, and Carol Propper. 2013. "Death by Market Power: Reform, Competition, and Patient Outcomes in the National Health Service." *American Economic Journal: Economic Policy* 5 (4): 134–66.

Gaynor, Martin, Carol Propper, and Stephan Seiler. 2016. "Free to Choose? Reform, Choice, and Consideration Sets in the English National Health Service." *American Economic Review* 106 (11): 3521–57.

Geruso, Michael, and Timothy Layton. 2020. "Upcoding: Evidence from Medicare on Squishy Risk Adjustment." *Journal of Political Economy* 128 (3): 984–1026.

Goldin, Jacob, Ithai Z. Lurie, and Janet McCubbin. 2021. "Health Insurance and Mortality: Experimental Evidence from Taxpayer Outreach." *Quarterly Journal of Economics* 136 (1): 1–49.

Gordon, Suzanne. 2019. "The VA Is Privatizing Veterans' Health Care While Launching a Campaign to Deny It." *American Prospect,* March 11. https://prospect.org/health/va-privatizing-veterans-health-care-launching-campaign-deny/.

Gottlieb, Daniel J., Weiping Zhou, Yunjie Song, Kathryn Gilman Andrews, Jonathan S. Skinner, and Jason M. Sutherland. 2010. "Prices Don't Drive Regional Medicare Spending Variations." *Health Affairs* 29 (3): 537–43.

Greaney, Thomas L. 2011. "Accountable Care Organizations: The Fork in the Road." *New England Journal of Medicine* 364 (1): 1.

Hart, Oliver, Andrei Shleifer, and Robert W. Vishny. 1997. "The Proper Scope of Government: Theory and an Application to Prisons." *Quarterly Journal of Economics* 112 (4): 1127–61.

Heckman, James J., and Edward J. Vytlacil. 2005. "Structural Equations, Treatment Effects, and Econometric Policy Evaluation." *Econometrica* 73 (3): 669–738.

Heckman, James J., and Edward J. Vytlacil. 2007. "Econometric Evaluation of Social Programs, Part II: Using the Marginal Treatment Effect to Organize Alternative Econometric Estimators to Evaluate Social Programs, and to Forecast their Effects in New Environments." In *Handbook of Econometrics*, Vol. 6B, edited by James J. Heckman and Edward E. Leamer, 4875–5143. Amsterdam: Elsevier.

Hendricks, Rickey. 1993. *A Model for National Health Care: The History of Kaiser Permanente Health and Medicine in American Society*. New Brunswick, NJ: Rutgers University Press.

Holmgren, A. Jay, Vaishali Patel, and Julia Adler-Milstein. 2017. "Progress in Interoperability: Measuring US Hospitals' Engagement in Sharing Patient Data." *Health Affairs* 36 (10): 1820–27.

Honore, Bo E., and Adriana Lleras-Muney. 2006. "Bounds in Competing Risks Models and the War on Cancer." *Econometrica* 74 (6): 1675–98.

House Veterans' Affairs Committee. 2018. *Maintaining Internal Systems and Strengthening Integrated Outside Networks* (*MISSION*) *Act of 2018*. Washington, DC: House Veterans' Affairs Committee.

Hull, Peter. 2020. "Estimating Hospital Quality with Quasi-Experimental Data." Unpublished.

Huskamp, Haiden A., and John K. Iglehart. 2016. "Mental Health and Substance-Use Reforms: Milestones Reached, Challenges Ahead." *New England Journal of Medicine* 375 (7): 688–95.

Hynes, Denise M., Samuel Edwards, Alex Hickok, Meike Niederhausen, Frances M. Weaver, Elizabeth Tarlov, Howard Gordon, et al. 2021. "Veterans' Use of Veterans Health Administration Primary Care in an Era of Expanding Choice." *Medical Care* 59 (6 Suppl 3): S292–S300.

Hynes, Denise M., Kristin Koelling, Kevin Stroupe, Noreen Arnold, Katherine Mallin, Min-Woong Sohn, Frances M. Weaver, Larry Manheim, and Linda Kok. 2007. "Veterans' Access to and Use of Medicare and Veterans Affairs Health Care." *Medical Care* 45 (3): 214–23.

Imbens, Guido W., and Joshua D. Angrist. 1994. "Identification and Estimation of Local Average Treatment Effects." *Econometrica* 62 (2): 467–75.

Imbens, Guido W., and Donald B. Rubin. 1997. "Estimating Outcome Distributions for Compliers in Instrumental Variables Models." *Review of Economic Studies* 64 (4): 555–74.

Jha, Ashish K., Catherine M. DesRoches, Eric G. Campbell, Karen Donelan, Sowmya R. Rao, Timothy G. Ferris, Alexandra Shields, Sara Rosenbaum, and David Blumenthal. 2009. "Use of Electronic Health Records in U.S. Hospitals." *New England Journal of Medicine* 360 (16): 1628–38.

Jha, Ashish K., Jonathan B. Perlin, Kenneth W. Kizer, and R. Adams Dudley. 2003. "Effect of the Transformation of the Veterans Affairs Health Care System on the Quality of Care." *New England Journal of Medicine* 348 (22): 2218–27.

Kefe, Josh. 2018. "Is the VA Being Privatized? Critics Say This Koch-Backed Group Is Fighting to Do Just That." *Newsweek*, April 5. https://www.newsweek.com/koch-brothers-backed-group-could-determine-future-va-870693.

Kelly, Elaine, and George Stoye. 2020. "The Impacts of Private Hospital Entry on the Public Market for Elective Care in England." *Journal of Health Economics* 73: 102353.

Kitagawa, Toru. 2015. "A Test for Instrument Validity." *Econometrica* 83 (5): 2043–63.

Kline, Patrick M., and Christopher R. Walters. 2019. "On Heckits, LATE, and Numerical Equivalence." *Econometrica* 87 (2): 677–96.

Kling, Jeffrey R. 2006. "Incarceration Length, Employment, and Earnings." *American Economic Review* 96 (3): 863–76.

Kolesar, Michal, Raj Chetty, John Friedman, Edward Glaeser, and Guido W. Imbens. 2015. "Identification and Inference with Many Invalid Instruments." *Journal of Business and Economic Statistics* 33 (4): 474–84.

Liu, Chuan-Fen, Adam Batten, Edwin S. Wong, Stephan D. Fihn, and Paul L. Hebert. 2018. "Fee-for-Service Medicare-Enrolled Elderly Veterans Are Increasingly Voting with Their Feet to Use More VA and Less Medicare, 2003-2014." *Health Services Research* 53 (S3): 5140–58.

McCarthy, Douglas, and David Blumenthal. 2006. "Stories from the Sharp End: Case Studies in Safety Improvement." *Milbank Quarterly* 84 (1): 165–200.

McWilliams, J. Michael, Laura A. Hatfield, Michael E. Chernew, Bruce E. Landon, and Aaron L. Schwartz. 2016. "Early Performance of Accountable Care Organizations in Medicare." *New England Journal of Medicine* 374 (24): 2357–66.

McWilliams, J. Michael, Bruce E. Landon, Michael E. Chernew, and Alan M. Zaslavsky. 2014. "Changes in Patients' Experiences in Medicare Accountable Care Organizations." *New England Journal of Medicine* 371 (18): 1715–24.

Miller, Amalia R., and Catherine E. Tucker. 2011. "Can Health Care Information Technology Save Babies?" *Journal of Political Economy* 119 (2): 289–324.

Miller, Sarah, Norman Johnson, and Laura R. Wherry. 2021. "Medicaid and Mortality: New Evidence from Linked Survey and Administrative Data." *Quarterly Journal of Economics* 136 (3): 1783–1829.

National Bureau of Economic Research. 1984–2023. "Diagnosis-Related Group Major Diagnostic Category Crosswalk." https://www.nber.org/research/data/diagnosis-related-group-major-diagnostic-category-crosswalk (accessed July 19, 2023).

Nevedal, Andrea L., Todd H. Wagner, Laura S. Ellerbe, Steven M. Asch, and Christopher J. Koenig. 2019. "A Qualitative Study of Primary Care Providers' Experiences with the Veterans Choice Program." *Journal of General Internal Medicine* 34 (4): 598–603.

Nuti, Sudhakar V., Li Qin, John S. Rumsfeld, Joseph S. Ross, Frederick A. Masoudi, Sharon-Lise T. Normand, Karthik Murugiah, Susannah M. Bernheim, Lisa G. Suter, and Harlan M. Krumholz. 2016. "Association of Admission to Veterans Affairs Hospitals vs Non-Veterans Affairs Hospitals with Mortality and Readmission Rates Among Older Men Hospitalized with Acute Myocardial Infarction, Heart Failure, or Pneumonia." *Journal of the American Medical Association* 315 (6): 582–92.

O'Hanlon, Claire, Christina Huang, Elizabeth Sloss, Rebecca Anhang Price, Peter Hussey, Carrie Farmer, and Courtney Gidengil. 2017. "Comparing VA and Non-VA Quality of Care: A Systematic Review." *Journal of General Internal Medicine* 32 (1): 105–21.

Olsen, Randall J. 1980. "A Least Squares Correction for Selectivity Bias." *Econometrica* 48 (7): 1815–20.

Perlman, Michael D. 1969. "One-Sided Testing Problems in Multivariate Analysis." *Annals of Mathematical Statistics* 40 (2): 549–67.

Phibbs, Ciaran S., Winifred J. Scott, Nicole E. Flores, and Paul G. Barnett. 2019. "HERC Outpatient Average Cost Dataset for VA Care." US Department of Veterans Affairs. http://vaww.herc.research.va.gov/include/page.asp?id=guidebook-outpatient-ac (accessed on April 22, 2020).

Ragone, Michael. 2012. "Evolution or Revolution: EMS Industry Faces Difficult Changes." *Journal of Emergency Medical Services* 37 (2): 34–39.

Reid, T. R. 2010. *The Healing of America: A Global Quest for Better, Cheaper, and Fairer Health Care.* London: Penguin Books.

Rein, Lisa, Philip Rucker, Emily Wax-Thibodeaux, and Josh Dawsey. 2018. "Trump Taps His Doctor to Replace Shulkin at VA, Choosing Personal Chemistry over Traditional Qualifications." *Washington Post*, March 29. https://www.washingtonpost.com/world/national-security/trump-ousts-veterans-affairs-chief-david-shulkin-in-administrations-latest-shake-up/2018/03/28/3c1da57e-2794-11e8-b79d-f3d931db7f68_story.html.

Rice, Thomas, Pauline Rosenau, Lynn Y. Unruh, and Andrew J. Barnes. 2013. "United States of America: Health System Review." *Health Systems in Transition* 15 (3): 1–431.

Rinne, Seppo T., Kirsten Resnick, Renda Soylemez Wiener, Steven R. Simon, and A. Rani Elwy. 2019. "VA Provider Perspectives on Coordinating COPD Care Across Health Systems." *Journal of General Internal Medicine* 34 (Suppl 1): 37–42.

Schwartz, Joel. 2000. "Harvesting and Long Term Exposure Effects in the Relation between Air Pollution and Mortality." *American Journal of Epidemiology* 151 (5): 440–48.

Shekelle, Paul G., Steven Asch, Peter Glassman, Sierra Matula, Amal Trivedi, and Isomi Miake-Lye. 2010. *Comparison of Quality of Care in VA and Non-VA Settings: A Systematic Review*. Washington, DC: Department of Veterans Affairs Technical Report.

Shulkin, David J. 2018. "Privatizing the V.A. Will Hurt Veterans." *New York Times*, March 28. https://www.nytimes.com/2018/03/28/opinion/shulkin-veterans-affairs-privatization.html.

Silver, David. 2020. "Haste or Waste? Peer Pressure and Productivity in the Emergency Department." Unpublished.

Skura, Barry. 2001. *Where Do 911 System Ambulances Take Their Patients? Differences between Voluntary Hospital Ambulances and Fire Department Ambulances*. New York, NY: Office Comptroller.

Sommers, Benjamin D., Atul A. Gawande, and Katherine Baicker. 2017. "Health Insurance Coverage and Health — What the Recent Evidence Tells Us." *New England Journal of Medicine* 377 (6): 586–93.

Starr, Paul. 2008. *The Social Transformation of American Medicine: The Rise of a Sovereign Profession and the Making of a Vast Industry.* New York, NY: Basic Books.

Stigler, George J. 1965. "The Economist and the State." *American Economic Review* 55 (1/2): 1–18.

Tanielian, Terri, and Lisa H. Jaycox, ed. 2008. *Invisible Wounds of War: Psychological and Cognitive Injuries, Their Consequences, and Services to Assist Recovery.* Santa Monica, CA: RAND Corporation.

Telser, Lester G. 1964. "Advertising and Competition." *Journal of Political Economy* 72 (6): 537–62.

Trivedi, Amal N., Sierra Matula, Isomi Miake-Lye, Peter A. Glassman, Paul Shekelle, and Steven Asch. 2011. "Systematic Review: Comparison of the Quality of Medical Care in Veterans Affairs and Non-Veterans Affairs Settings." *Medical Care* 49 (1): 76–88.

Trombley, Matthew J., Betty Fout, Sasha Brodsky, J. Michael McWilliams, David J. Nyweide, and Brant Morefield. 2019. "Early Effects of an Accountable Care Organization Model for Underserved Areas." *New England Journal of Medicine* 381 (6): 543–51.

US Government Accountability Office. 2015. *VA Mental Health: Clearer Guidance on Access Policies and Wait-Time Data Needed*. Washington, DC: U.S. Government Accountability Office.

Veterans Access, Choice, and Accountability Act of 2014, Pub. L. No. 113-146, (2014).

Veterans Millennium Health Care and Benefits Act, Pub. L. No. 106-117, 113 Stat. 224 (1999).

Wagner, Todd H., Shuo Chen, and Paul G. Barnett. 2003. "Using Average Cost Methods to Estimate Encounter-Level Costs for Medical-Surgical Stays in the VA." *Medical Care Research and Review* 60 (3): 15S–36S.

Wasserman, Jeffrey, Jeanne S. Ringel, Karen A. Ricci, Jesse D. Malkin, Barbara O. Wynn, Jack Zwanziger, Sydne Newberry, et al. 2005. *Analyzing and Influencing How the VA Allocates Its Health Care Dollars*. Santa Monica, CA: RAND Corporation.

Wolak, Frank A. 1987. "An Exact Test for Multiple Inequality and Equality Constraints in the Linear Regression Model." *Journal of the American Statistical Association* 82 (399): 782–93.

Woodworth, Lindsey. 2020. "Swamped: Emergency Department Crowding and Patient Mortality." *Journal of Health Economics* 70: 102279.

Wooldridge, Jeffrey M. 2015. "Control Function Methods in Applied Econometrics." *Journal of Human Resources* 50 (2): 420–45.

Wright, Steven M., Laura A. Petersen, Rebecca P. Lamkin, and Jennifer Daley. 1999. "Increasing Use of Medicare Services by Veterans with Acute Myocardial Infarction." *Medical Care* 37 (6): 529–37.