

NBER WORKING PAPER SERIES

THE PRODUCTIVITY OF PROFESSIONS:
EVIDENCE FROM THE EMERGENCY DEPARTMENT

David C. Chan Jr
Yiqun Chen

Working Paper 30608
<http://www.nber.org/papers/w30608>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2022, Revised August 2026

We are grateful to Diane Alexander, Ricardo Alonso, Benjamin Arold, Sandy Black, Kate Bundorf, Marika Cabral, David Card, Stuart Craig, Janet Currie, Shooshan Danagouliau, Qing Gong, Joshua Gottlieb, Mitch Hoffman, Tom Hubbard, Bapu Jena, Amanda Kowalski, Brad Larsen, Darren Lubotsky, Bentley MacLeod, Neale Mahoney, Benjamin McMichael, David Molitor, Jessica Monnet, Ciaran Phibbs, Maria Polyakova, Julian Reif, Michael Richards, Steve Rivkin, Evan Rose, Susan Schmitt, Molly Schnell, Brad Shapiro, Isaac Sorkin, Chris Walters, and many seminar and conference participants for helpful comments and suggestions. Sam Bock, Noah Boden-Gologorsky, Akriti Dureja, Jesse Kozler, Matthew Merrigan, Francis Peng, Jonatas Prates, Aadit Shah, Kemin Wang, Justine Weng, and Saam Zahedian provided excellent research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by David C. Chan Jr and Yiqun Chen. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Productivity of Professions: Evidence from the Emergency Department

David C. Chan Jr and Yiqun Chen

NBER Working Paper No. 30608

October 2022, Revised August 2026

JEL No. I11, I18, J24, J44, M53

ABSTRACT

This paper studies the productivity of nurse practitioners (NPs) and physicians, two professions performing overlapping tasks but with starkly different backgrounds, training, and pay. Using quasi-experimental variation in patient assignment to NPs versus physicians in Veterans Health Administration emergency departments, we find that, on average, NPs use more resources and exhibit a higher 30-day preventable hospitalization rate than physicians. However, the NP-physician performance difference varies by case complexity and severity. Importantly, even larger productivity variation exists within each profession, leading to substantial overlap between the productivity distributions of the two professions; NPs outperform physicians in 38 percent of random pairs.

David C. Chan Jr

University of California, Berkeley

Haas School of Business

and NBER

david.c.chan@berkeley.edu

Yiqun Chen

University of Illinois Chicago

Department of Economics

and NBER

yqchen@uic.edu

1 Introduction

Professions play a key role in determining the division of labor and the returns to skilled work (Abbott 2014). In selecting and training future members, professional groups may both restrict the supply of professionals and impact their quality. Professional groups may also act as special interest groups for their members, lobbying policymakers and negotiating with payers. Large differences in pay between professional groups may, therefore, reflect differences in worker productivity or rents from restricted supply or privileged arrangements (Freidson 1974; Shapiro 1986).

Evidence comparing the productivity of distinct classes of professionals performing overlapping tasks remains scant. Professional groups, by nature, act to exclude outsiders from their “jurisdictions” (Abbott 2014). While the medical profession provides a well-documented case study of historical exclusion, it now provides an opportunity for study. Recent decades have witnessed an increased demand for health care outstripping the supply of physicians and the rise of NPs seeking to perform some of the same tasks that physicians do. The number of NPs has reached about one-third of the number of physicians in the US (Bureau of Labor Statistics 2021a; 2021b), and various states have responded to the shortfall in physicians by liberalizing NP “scope of practice” to perform tasks traditionally performed by physicians (e.g., McMichael and Markowitz 2022). While policy discussions of NPs have largely highlighted their role in primary care, NPs practice in a wide range of areas. Around one-third of NPs work in primary care; the remaining work in various specialty settings.¹

In this paper, we exploit a quasi-experiment in the Veterans Health Administration (VHA) to study the productivity of NPs versus physicians. In December 2016, the VHA granted full practice authority to NPs, allowing them to practice without physician supervision. We leverage quasi-experimental variation in the availability of physician and NP providers in the emergency department (ED). Nationwide, the share of ED visits seen by NPs reached 13 percent in 2019 (Cairns and Kang 2022); 4 percent of NPs work in emergency care, close to the share of physicians specializing in emergency medicine.² In a sample of 1.1 million ED visits, our approach compares patients arriving at the same ED and during similar times (i.e., the year, month, day of the week, and hour of the day) that differ in the number of NPs on duty. We show that the number of on-duty physicians declines with the number of on-duty NPs, and the number of NPs on duty strongly predicts whether an arriving patient will be assigned to an NP versus a physician. Under the plausible assumption that patients arrive quasi-randomly within cells of ED stations and time categories, this

1. See Measure 3.2 of Milbank Memorial Fund (2024) and Table 5-2 of Health Resources and Services Administration (2010).

2. According to the National Sample Survey of Registered Nurses (NSSRN 2022), 4 percent of NPs work in emergency care; according to the Association of American Medical Colleges (AAMC 2021), 5 percent of physicians specialize in emergency medicine.

instrumental variables (IV) design allows us to study the effect of NPs on patient resource utilization and health outcomes.

The results show that, on average in the ED setting, NPs use more resources during the visit and exhibit a higher 30-day preventable hospitalization rate than physicians. On average, NPs increase patient length of stay by 11 percent and raise the cost of ED care by 7 percent. We do not find NP effects on inpatient admissions in the overall sample, though NPs raise admissions for high-risk cases. For downstream patient outcomes, NPs on average raise 30-day preventable hospitalizations, a frequently used quality-of-care measure, by 20 percent relative to the mean.³ We do not find statistically significant NP effects on 30-day mortality in the overall sample (95 percent confidence interval: -0.34 to 0.11 percentage points). In contrast to IV estimates, ordinary least squares (OLS) estimates for resource use and 30-day preventable hospitalizations are negative in sign, consistent with the descriptive evidence that NPs treat healthier patients than physicians do.

Under various analytical lenses, we uncover mechanisms behind and responses to the average performance difference between NPs and physicians in the ED. First, we find suggestive evidence that experience matters: The NP-physician gap in some outcomes narrows among providers who have seen more prior patients, both in general and for the diagnosis in question. This indicates that differences in training may play some role in the performance differences between NPs and physicians. Second, we document differences in clinical decision-making. Compared to physicians, NPs are likelier to obtain external information from radiology tests and formal consults. NPs also exhibit prescription patterns consistent with responses to lower skill (Chan, Gentzkow, and Yu 2022): Relative to physicians, NPs are *less* likely to prescribe opioids, which have higher health risks if incorrectly prescribed, but they are *more* likely to prescribe antibiotics, which have higher health risks if incorrectly not prescribed. Despite lower opioid prescription rates, NPs obtain similar downstream rates of opioid use disorder, and despite higher antibiotic prescription rates, NPs exhibit higher rates of subsequent ED visits with infection. Third, we examine heterogeneity in the NP-physician gap by patient condition complexity and severity. The NP-physician gap in patient length of stay, ED cost, and the probability of inpatient admission is smaller for patients with fewer comorbidities and lower severity. Finally, case assignment seems to respond to NPs' relative disadvantage in treating complex cases and absolute disadvantage: On average, NPs are assigned healthier patients of those available for assignment; NPs are also assigned a (modestly) smaller share of patients when the ED is less busy.

The heterogeneity of our results by case complexity and severity complements the earlier literature on the productivity of NPs relative to physicians. Most of the previous studies focus on the primary care setting.

3. Preventable hospitalizations are defined using the standardized algorithm developed by the Agency for Healthcare Research and Quality and relied upon by researchers and policymakers to measure quality (e.g., New Jersey Department of Health 2011; Nyweide et al. 2013; California Health and Human Services Agency 2025; Currie and Zhang 2025).

With the caveat that those studies feature small samples with possibly underpowered estimates or empirical strategies with limited causal interpretations, our results of smaller NP-physician performance differences in the ED for less complex and less severe cases may align with the earlier literature that indicates greater parity between NPs and physicians in primary care, providing a policy-relevant message that suggests roles for both the NP and physician professions in health care delivery.⁴

We undertake several analyses to assess the validity of our IV quasi-experiment. We show that, conditional on the baseline controls, a broad range of patient characteristics that predict outcomes is well balanced across values of our instrument. Relatedly, our IV estimates are remarkably stable, regardless of the inclusion of a wide set of patient covariates; in contrast, OLS estimates under the full set of covariates are less than half the magnitude (but still opposite-signed relative to IV estimates) of those when only baseline covariates are included. We also assess the validity of the exclusion restriction, that the number of NPs on duty is not correlated with other factors that could drive care delivery or patient outcomes. Specifically, we show that our instrument is conditionally unrelated to the characteristics of on-duty physicians and NPs. We further examine potential spillovers between NPs and physicians and find no evidence suggesting such spillovers.⁵ Finally, we show that our results are robust to controlling for a series of other factors that may vary with the instrument, including patient volume, total staffing, and patient wait times.

We perform counterfactual analyses to understand the magnitude of the average NP-physician productivity difference, relative to differences between physician and NP wages. We find that compared to the counterfactual scenario of no NPs, allocating one quarter of ED patients to NPs increases non-wage spending by \$197 million per year to the VHA. Accounting for lower NP wages that are half that of physicians' wages, it again implies a net cost of \$129 million per year. However, the cost implications of NPs decline with case complexity. Based on IV estimates, using NPs to staff the least-complex quarter of cases (based

4. An older medical literature has raised the question of the performance of NPs relative to physicians, but the generally small numbers of providers and other features of these earlier studies limit inference on systematic differences between the two classes of providers. See Laurant et al. (2005) for a systematic review of this literature. The literature features small randomized trials with null results, sometimes comparing a single-digit number of physicians with a single-digit number of NPs. The studies tend to focus on primary care settings and usually have short follow-up times that may be insufficient to detect meaningful effects. According to Laurant et al. (2005), the null findings of this literature "should be viewed with caution given that only one study was powered to assess equivalence of care, many studies had methodological limitations, and patient follow-up was generally 12 months or less." Some recent studies leverage larger-scale data, yet patient selection between NPs and physicians limits the causal inference (e.g., Perloff, DesRoches, and Buerhaus 2016; Lutfiyya et al. 2017; Liu et al. 2020). A notable exception is Currie and Zhang (2025), which finds that non-physician providers—physician assistants and NPs combined—may not perform worse than physicians in primary care in the VHA. Yet the paper's study period, January 2005 to February 2017, is before physician assistants were legally allowed to practice independently at the VHA and mostly before the VHA implementation of NP full practice authority. Thus, non-physician providers in the paper's setting were largely mandated to practice under physician supervision or oversight.

5. Specifically, we examine whether NP presence may affect physician performance. E.g., NPs asking physicians for assistance could slow down physicians. We thus examine whether physicians' treatment outcomes change with the presence of NPs; we find no such evidence. We also examine whether on-duty physician quality impacts NP performance; we find that the outcomes of patients treated by NPs are unrelated to the value-added of the physicians on duty.

on Elixhauser comorbidities) results in net spending that is only about one-fifth of the amount for assigning a quarter of cases at the average complexity.

Finally, we examine variation in productivity across providers within the professional classes of NPs and physicians. To arrive at provider-specific measures of productivity, we estimate a just-identified IV model, in which we instrument patient assignment to specific providers by indicators for on-duty providers. Using a method developed by Efron (2016) and adapted by Kline, Rose, and Walters (2022), we deconvolve the estimates of provider-specific productivity into underlying productivity distributions for each of the two professional classes. We find wide variation in productivity within professions and substantial overlap between the productivity distributions of NPs and physicians. The probability that a randomly chosen NP is more productive than a randomly chosen physician can be as large as 38 percent. Within each professional class, the productivity distribution implies greater medical spending of around \$900,000 per year under a provider at the 25th percentile of productivity than under a provider at the 75th percentile—about three times the mean annual spending difference between NPs and physicians. That is, despite the stark differences in training and selection, there is a sizable productivity overlap between NPs and physicians. While training and selection are relatively uniform within professions, productivity variation within professions is even larger than that between professions.

We extend our distributional analysis of productivity to examine the complexity of cases assigned to and the wages paid to individual providers. While productivity differences are much larger within each profession than the average difference between NPs and physicians, within each profession, a provider's productivity shows a limited relationship with her wages or the complexity of her assigned cases. This starkly contrasts with differences in case assignment and wages between professions: EDs assign noticeably less complex cases to NPs—enough to reverse OLS estimates for key outcomes relative to the corresponding IV estimates—and pay NPs half the wages of physicians. In other words, a provider's productivity is far from a sufficient statistic explaining her pay or determining her case assignment; professional class is a much stronger predictor of both wages and assigned case complexity.

Our findings contribute to several strands of literature. First, given the dramatic rise in the supply of NPs to meet the growing demand for health care, policy debates have arisen around the NP provision of care. A recent body of research has examined the effects of liberalizing state “scope of practice laws” for NPs.⁶ By design, these papers study the general-equilibrium impacts of both allowing NPs greater scope to practice

6. The findings of this literature are varied and somewhat mixed. Perry (2009), Kleiner et al. (2016), and Dillender et al. (2022) find that these laws affect NP earnings. Stange (2014) finds a minimal impact of greater NP supply on utilization, access, or prices but perhaps a moderate impact on primary care utilization. Yet Traczynski and Udalova (2018) find increases in utilization and some evidence of increased quality, and Alexander and Schnell (2019) find evidence of better access and outcomes in mental health.

and increasing the supply of providers; results will depend on how labor is reallocated between professions. Our study sheds light on the effect of assigning patients to NPs versus physicians. The widely heterogeneous complexity and severity of cases in our setting also provide a broader picture of the productivity of NPs and physicians. The heterogeneity of our results by patient type suggests returns to optimal case assignment for NPs. Also, the substantial overlap between the productivity distributions of NPs and physicians, even in the ED, may generate important policy implications supporting the use of NPs in less complex settings.

Strikingly, we find that productivity variation within the NP and physician professions is much larger than the average productivity difference between the two professions. This finding puts the prior literature on practice variation (e.g., Epstein and Nicholson 2009; Gowrisankaran, Joiner, and Léger 2023) in context: In our setting, the interquartile range in productivity within professions is several-fold larger than the overall productivity difference between two professions with starkly different selection and training processes. In contrast to the assignment of cases between NPs and physicians, in which less complex cases are, on average, assigned to NPs, we find essentially no matching between case complexity and provider productivity within professions. This suggests substantial informational and organizational barriers to such “skill-task matching” within professions, though professional boundaries may provide a mechanism for this matching between professions (Acemoglu and Autor 2011).

Second, our research relates to the widespread practice of occupational licensing, affecting about a third of all jobs in the US (Kleiner and Krueger 2013). The existing literature suggests that occupational licensing increases the earnings of licensed workers (Kleiner and Krueger 2013; Kleiner et al. 2016; Farronato et al. 2020) but provides little evidence on whether higher earnings arise from restricting the supply of workers or from improving the quality of their work in modern settings (see Farronato et al. 2020 for a review).⁷ Studies in this literature compare workers within professions (along the margin of occupational licensing), while ours compares two competing professions with important differences beyond licensing.

A third related literature is concerned with worker human capital and productivity. These issues have received growing attention in health care (e.g., Doyle, Ewer, and Wagner 2010; Currie and MacLeod 2017; 2020; Chen 2021; Chan, Gentzkow, and Yu 2022) and more broadly (e.g., Gennaioli et al. 2013; Chetty, Friedman, and Rockoff 2014). Our study contributes to this literature by comparing the productivity of distinct professional classes with starkly different selection and training. In the rich ED setting, we also uncover key mechanisms connecting productivity to human capital along dimensions of experience,

7. Two studies of an earlier, unregulated environment of midwifery, near the beginning of the 20th century, demonstrate meaningful reductions in maternal and infant mortality with the initial implementation of occupational licensing (Lazuka 2018; Anderson et al. 2020). In the education literature, Carrell and West (2010) find that (mostly tenure-track) professors with doctorate degrees produce better long-term student outcomes than (mostly adjunct) professors without such degrees.

information-gathering, decision-making, and case complexity. Our results suggest that professional selection and training may give rise to important productivity implications, though these implications may be small relative to the variation within professions.

Fourth, a broad set of questions concerns the distribution of wealth in society across occupations and strata of educational attainment. In recent decades, societies have witnessed an increased concentration of wealth in occupations associated with high human capital (Smith et al. 2019). Training to reach the highest levels of income has become increasingly competitive among the upper class, while the middle and lower classes are increasingly left behind, characterizing a “meritocracy trap” (Markovits 2020). Interestingly, our results suggest a productivity difference between professions at least as large as their wage difference, at least in our resource-intensive and information-dependent setting within health care. Yet, the lack of relationship between worker-specific productivity and wages within professions suggests frictions in observing or contracting by actual productivity across similar workers (e.g., Acemoglu and Pischke 1998). Entering a profession may represent a costly and imperfect way for workers to distinguish themselves.

The remainder of this paper is organized as follows. Section 2 describes the institutional setting and the data. Section 3 describes our empirical approach and provides evidence for its validity. Section 4 presents the average performance difference between NPs and physicians and evidence on mechanisms and responses. Section 5 presents analyses on policy-relevant counterfactual scenarios. Section 6 reports distributions of productivity, case assignment, and wages within professions. Section 7 concludes.

2 Background and Data

2.1 NPs and Physicians in the US

To understand the physician and NP professions in the US, it is instructive to consider their distinct origins in the American context. According to the landmark work by Starr (1982), “among the professions, medicine is both the paradigmatic and the exceptional case: paradigmatic in the sense that other professions emulate its example; exceptional in that none have been able to achieve its singular degree of economic power and cultural authority.” Aided by scientific advances and demographic shifts, the US medical profession in the early twentieth century captured authority by standardizing education and licensing toward a scientific orientation, in the process excluding a large swath of practitioners (Brown 1979; Larson 1979).⁸

The “professionalization” of American nursing began in the 1870s, with the establishment of the first

8. Much of this transformation centered around the Flexner Report in 1910, which strongly advocated for a scientific orientation of medical education and the exclusion of alternative practices (Flexner 1910; Beck 2004). Following the report, more than half of medical schools closed or consolidated (Patel and Rushefsky 2004).

three nursing schools in the US. In contrast to the scientific orientation of the medical profession, the driving force behind the nursing profession was to install (female) staff in hospitals to improve hygiene and cleanliness (Ashley 1976). NPs emerged from the nursing tradition in the 1960s, in the setting of increased specialization in medicine, worsening access to care, and new federal funding from Medicare and Medicaid to increase the training of providers (Fairman 2009; Hallett 2016). Over the next few decades, pressures to contain costs and increase throughput further expanded the boundaries of NP practice in a variety of settings (e.g., primary care, emergency care) (Fairman 2009; Kleinpell, Cook, and Padden 2018). While NPs have historically focused on the provision of primary care, which is still the largest area employing NPs, NPs now work in various specialty settings. Only around one-third of NPs worked in primary care between 2016 and 2021, approximately the share of physicians working in primary care (Milbank Memorial Fund 2024). Anecdotal evidence suggests that some of the same pressures that have lured physicians into specialty care have also begun to take hold among NPs (Andrews and Beard 2024).

Present-day NPs and physicians remain starkly different regarding training and income, despite performing overlapping tasks in many settings. Admission rates to medical schools are only about one-tenth admission rates to graduate nursing programs,⁹ and the number of years of training for physicians is around twice that for NPs.¹⁰ Undergoing a highly selective process and long training periods, physicians constitute the most common profession in the top percentile of the income distribution (Gottlieb et al. 2020); the income of NPs is roughly half that of physicians (Bureau of Labor Statistics 2021a; 2021b).

2.2 VHA and ED Setting

In December 2016, the VHA granted full practice authority to NPs. The policy enables NPs to practice without requiring physician supervision at VHA facilities. NPs can treat patients independently, regardless of state restrictions that would otherwise limit NPs' practice authority.¹¹

Several features make the ED a setting well suited to studying provider productivity. First, each patient visit is generally assigned to a single ED provider. Such independence allows us to attribute patient outcomes

9. We searched <https://www.petersons.com/graduate-schools.aspx> for admission rates to graduate nursing programs and to medical schools in the following universities with both graduate nursing programs and medical schools: Columbia University, Duke University, Emory University, Johns Hopkins University, the University of North Carolina, the University of Pennsylvania, and the University of Washington. Admission rates ranged from 25 to 63 percent for graduate nursing programs and from 3 to 7 percent for medical schools.

10. According to AAMC (2020), physicians must complete a four-year undergraduate degree, a four-year Doctor of Medicine degree, and three to seven years of residency training. Some physicians further undergo one to three years of fellowship training after residency. According to the American Association of Nurse Practitioners (AANP 2020), NPs must complete a four-year Bachelor of Science in Nursing degree and may choose between one to two years in a Master of Science in Nursing degree or three to four years in a Doctor of Nursing Practice degree. There is no residency or fellowship training requirement for NPs.

11. As of 2021, about half of the states in the US had not granted NPs full practice authority (AANP 2021).

to individual ED providers. Second, patient flow in the ED is highly unpredictable, while provider schedules are typically set well in advance. Variation in NP availability is thus unrelated to the set of patients arriving. Third, the ED is an important setting that uses the NP workforce. Across the nation in 2019, 13 percent of ED visits were seen by NPs (Cairns and Kang 2022).¹² In the VHA, the share of ED visits seen by NPs was 11 percent in 2019—about half the share of visits seen by NPs in primary care at around 20 percent in the VHA (Morgan et al. 2012).¹³ Nationwide, around 4 percent of NPs practice in emergency care, close to the share of physicians specializing in emergency medicine at 5 percent (AAMC 2021; NSSRN 2022). Fourth, patients present at the ED with a wide spectrum of conditions, ranging in complexity and severity, which provides an opportunity to investigate productivity across a range of tasks.

In Appendix Table A.1, we report the characteristics of NPs and physicians working at the VHA and non-VHA EDs. NPs in VHA EDs are representative of their non-VHA counterparts in female share (about 80 percent), while they appear to be more experienced (as indicated by age, 51 versus 43 years old). Among ED physicians, those practicing at the VHA are slightly more likely to be female than those outside of the VHA (34 versus 27 percent), but the two groups have a similar average age (48 versus 46 years old). Appendix Table A.1 also compares NPs practicing at the ED and the overall NP workforce, showing that ED NPs are representative of the NP workforce in age and gender.

2.3 Data

Our main data are administrative health records from the VHA, the largest integrated health care system in the US, serving more than nine million veterans. For each ED visit, the data record the type of provider treating the patient (i.e., NP or physician), resource utilization, and patient outcomes. The data also contain detailed information on patient characteristics (e.g., demographics, comorbidities, vital signs) and provider characteristics (e.g., birthdate, gender). The large sample size, as well as the availability of detailed information (e.g., time stamps for measuring ED length of stay, provider orders), enable a comprehensive analysis of NP- and physician-provided care. Finally, we access detailed payroll records for providers to impute wages.

Sample Construction.—We restrict our analysis sample in the following ways. First, we restrict the sample to ED visits between January 2017 and January 2020, i.e., after full practice authority was granted to NPs at the VHA and before the onset of the COVID pandemic in the US. Second, we include only cases arriving during the daytime (8 a.m. to 6 p.m.), because the data show that few NPs take evening or night

12. About 6 percent of ED visits nationwide were seen by NPs without physicians; 7 percent were seen by NPs with physicians.

13. We estimate the share of ED visits seen by NPs from the VHA data, which is described in the following section.

shifts.¹⁴ Third, we focus on visits to VHA EDs in months when the ED used NPs to treat patients. Local VHA facilities varied in whether they used NPs in the ED.¹⁵ Fourth, to examine a single margin between NPs and physicians, we exclude EDs that used non-physician providers other than NPs (mainly physician assistants).¹⁶ Finally, we drop a small number of cases with missing age or gender or aged under 20 or above 99 years. Appendix Table A.2 reports the sample size at each step. The final sample contains 1.1 million cases, over 44 EDs.¹⁷

Outcome Variables.—To measure medical resource use, we include two primary outcomes that are frequently used in the ED setting: (i) the patient length of stay (i.e., the time between patient assignment to the provider and patient discharge) and (ii) the cost of care during the ED visit (excluding costs due to a resulting hospital admission, measured separately next).¹⁸ We also include hospital admission as a resource use measure, indicating the provider’s decision to admit the patient for inpatient care during the ED visit. To measure the quality of care, we examine two prominent patient outcomes: We use linked death records to construct indicators of patient 30-day mortality and use linked inpatient data to construct indicators of 30-day preventable hospitalization as defined by the Agency for Healthcare Research and Quality (2021). We measure 30-day preventable hospitalization after ED discharge; thus, it is beyond the decision of the provider in the index ED visit.

In examining mechanisms behind the effect of NPs, we include the following sets of outcomes: (i) whether the provider orders consults (which are typically from specialists outside of the ED); (ii) whether the provider orders CT scans and X-rays, two primary diagnostics in the ED; (iii) whether the provider prescribes opioids and antibiotics, two major classes of drugs whose clinical indications for appropriate use are often unclear and can require skill to discern (e.g., Fleming-Dutra et al. 2016; Neuman, Bateman, and Wunsch 2019); and (iv) whether the patient has return visits with infections and subsequent opioid use disorder.

14. One possible explanation is that, since patient volumes on average are much lower in the evening/night than in the daytime (e.g., the average number of cases arriving per hour is 3.6 between 8 a.m. and 6 p.m. versus only 0.9 outside of 8 a.m.-6 p.m.), EDs in our sample often staff only one provider per shift in the evening/night. Our interview with VHA ED leadership revealed that EDs are required to have at least one physician on duty at all times.

15. We define a VHA ED as using NPs in a month if it has at least 15 cases treated by NPs in the month. The sample size changes only slightly when we use alternative thresholds: e.g., the sample size is 1.13 and 1.10 million when using a threshold of 10 and 20, respectively, compared to 1.12 million based on a threshold of 15.

16. In addition, unlike NPs, physician assistants had not been allowed to practice independently at the VHA as of 2022 nor in any state as of mid-2019. In Section 4.6, we expand the sample to cases in all EDs that use NPs despite the use of physician assistants but exclude cases in ED-day cells with physician assistants to focus on the margin between NPs and physicians. This analysis expands the sample from 1.1 million cases across 44 EDs seen by 1,348 physicians and 156 NPs in our baseline analysis to 2.2 million cases across 110 EDs seen by 3,499 physicians and 491 NPs. We find similar results in this expanded sample.

17. Appendix Table A.3 shows that EDs included in our sample are a representative set of all VHA EDs, as measured by ED, provider, and patient characteristics. The EDs in our sample include those in large metropolitan areas (e.g., New York City, Los Angeles, the Bay Area) as well as those in smaller areas (e.g., Detroit, Michigan; Madison, Wisconsin).

18. For length of stay, we use detailed time-stamped data on patient assignment and discharge. For the cost of ED care, we use cost accounting by the VHA, which measures resource utilization.

Descriptive Statistics.—Table 1 summarizes the characteristics of the cases included in our analysis. In addition to demographics, we measure patient comorbidities as Elixhauser indices (Elixhauser et al. 1998), which are 31 indicators for comorbidities (e.g., cancer, diabetes) based on patient medical histories in the prior 365 days. We also report the average length of stay, average ED spending (inflation-adjusted to 2020 dollars), and 30-day preventable hospitalization rate. Column 1 shows the characteristics of the overall sample. Columns 2 and 3 compare cases treated by NPs with those treated by physicians. Along several dimensions, cases treated by NPs are healthier than those treated by physicians: Cases treated by NPs are younger (60.7 versus 62.5 years old), have fewer Elixhauser comorbidities (3.2 versus 3.7), and have fewer outpatient visits and fewer inpatient stays in the prior 365 days (5.7 versus 6.4 outpatient visits and 0.4 versus 0.7 inpatient stays). Consistent with selection, cases treated by NPs appear to have better outcomes: They have a shorter average length of stay (120 versus 175 minutes), a lower average ED cost (\$813 versus \$978), and a lower 30-day preventable hospitalization rate (0.8 versus 1.4 percentage points).

3 Empirical Strategy

An ideal experiment to assess the effect of being treated by NPs would randomly assign cases to NPs and physicians. Lacking random assignment, we use a quasi-experimental approach: We leverage plausibly exogenous variation in the availability of NPs on duty to instrument for whether an NP or a physician treats a case. In this section, we begin by describing our instrumental variables (IV) approach. We then show evidence supporting the validity of our identification strategy.

3.1 Specification

Our empirical specification is a two-stage least squares (2SLS) model that takes the following form:

$$y_i = \delta \text{NP}_i + \mathbf{T}_i \eta + \mathbf{X}_i \beta + \varepsilon_i, \quad (1)$$

$$\text{NP}_i = \lambda Z_i + \mathbf{T}_i \zeta + \mathbf{X}_i \gamma + v_i, \quad (2)$$

where i denotes a case, y_i is the outcome of interest, and NP_i indicates whether case i is treated by an NP. We use Z_i to denote the instrument, i.e., the number of NPs on duty between 8 a.m. and 6 p.m. (our analysis time window) at the ED on the day case i visits.¹⁹ The parameter of interest is δ , which represents a local average

19. Since the data do not include direct information on provider scheduling, we measure Z_i as the number of NPs treating cases arriving during the analysis time window in the ED-day cell of case i 's visit. We count an NP as on duty if she is observed treating at least two cases between 8 a.m. and 6 p.m. in the ED-day cell. The main analysis includes the index case in calculating Z_i . In Section 4.6, we show the robustness of our estimates to alternative measurements of Z_i , including (i) one that includes NPs with

treatment effect (LATE): the average causal effect among cases that would have been assigned to a different class of provider under a different number of NPs on duty. In Section 4.6, we include alternative instruments, including the share of on-duty providers who are NPs, the share of cases assigned to NPs (leaving out the index case), and an indicator for any NPs on duty in the ED-day cell.

The vector $\mathbf{T}_i$ encodes interactions between indicators for the ED and indicators for time categories of the patient’s arrival. Specifically, $\mathbf{T}_i$ contains ED-by-year, ED-by-month, ED-by-day-of-the-week, and ED-by-hour-of-the-day indicators. We condition on $\mathbf{T}_i$ to allow for the sorting of NPs across shift types (e.g., weekdays versus weekends) and EDs, where patient characteristics and ED conditions may be systematically different. Controlling for ED-by-time-category indicators captures these potential systematic differences.²⁰

As robustness checks, our specification also includes a vector of patient covariates $\mathbf{X}_i$, including indicators for five-year age bins, marital status, gender, and race (white, Black, and Asian/Pacific Islander, with other racial categories omitted as the reference group); indicators for 31 Elixhauser comorbidities; prior health care use (the number of outpatient visits and the number of inpatient stays in VHA facilities in the prior 365 days); vital signs (pulse, respiratory rate, blood oxygen level, pain level, body temperature, an indicator for fever, systolic blood pressure, and diastolic blood pressure); and indicators for the three-digit ICD-10 code of the primary diagnosis of the visit.²¹ For each patient covariate with missing values, we add an indicator for missing values and replace the missing values with zero. Finally, ε_i and v_i are error terms. We cluster standard errors by provider. We also consider alternative clustering approaches in Section 4.6.

3.2 Identification

To interpret δ as the LATE of being treated by NPs, our IV approach requires four identifying assumptions: relevance, conditional independence, exclusion, and monotonicity. In this section, we summarize empirical evidence supporting the validity of the identifying assumptions.

Relevance.—Figure 1 shows the first stage of our IV model, controlling for the baseline controls $\mathbf{T}_i$. Panel A shows that the number of physicians on duty declines linearly with the number of NPs on duty. Consequently, Panel B shows that a patient’s probability of being treated by an NP increases with the number of NPs on duty: One more NP on duty increases her probability of being treated by an NP by 18.6 percent.

only one case in a shift and (ii) one that leaves out the index case in defining whether an NP is on duty.

20. While fixed effects for ED-by-hour-of-the-day interactions are not necessary to condition on to yield quasi-random variation in the instrument (since the instrument varies at the day instead of hour level), we include them for statistical precision of estimates. This inclusion does not meaningfully affect the statistical significance of our results.

21. Since our study period is from January 2017 to January 2020, all disease diagnoses are coded in ICD-10 in the data. A potential question is whether the three-digit diagnoses are endogenous to being treated by NPs. Yet, as shown in Section 4 below, our estimates are remarkably stable regardless of whether we control for three-digit diagnosis indicators. In Appendix A.1, we also show that NPs and physicians are similar in their coding of three-digit diagnoses.

The increase is highly significant (with an F -statistic of 149.2, even conditioning on $\mathbf{T}_i$) and is close to linear.

To provide context, Appendix Figure A.1 presents a histogram of the number of NPs on duty across ED-day cells. The figure reveals a fair spread in the number of NPs on duty: 38.1, 47.2, and 11.3 percent of ED-day cells have zero, one, and two NP(s) on duty, respectively; 3.4 percent of ED-day cells have more than two NPs on duty. An important question is what drives the variation in the number of NPs on duty. As shown in Appendix Figure A.2, NPs are less likely to work on weekends: 77 percent of weekday ED-day cells versus 24 percent of weekend ED-day cells have NPs on duty. However, conditional on the day of the week and other time categories (year and month), we still observe a large variation in the number of NPs on duty within EDs across days (standard deviation of 0.49).

Interviews and surveys with VHA ED providers and administrators point to sources of random variation in scheduling:²² As EDs have long open hours (typically 24/7), EDs rotate providers on duty. Providers' on-duty schedules are generally set months in advance, restricting the possibility of arranging on-duty providers based on specific conditions in the ED on the day of occurrence, e.g., patient volumes or diseases of patients arriving. Each ED has a staffing model that pre-specifies the staffing level for each type of shift defined by the time categories (matching $\mathbf{T}_i$) to meet patient care needs predicted by the time categories. In general, conditional on $\mathbf{T}_i$, there is no systematic variation in the characteristics of the staff scheduled to be on duty.²³ It is also extremely rare that EDs call in providers who are not scheduled to be on duty to care for patients.

In the empirical evidence below, we show that patient characteristics, as well as on-duty physicians' and on-duty NPs' characteristics, are well balanced across the number of NPs on duty (conditional on $\mathbf{T}_i$), consistent with quasi-random variation in the instrument. Otherwise, e.g., if the variation is driven by patient or staffing shocks, we would observe systematic changes in patient and/or on-duty physicians' and NPs' characteristics with the instrument.

In Appendix Figure A.3, we decompose the variance of NP staffing. The figure shows that $\mathbf{T}_i$ explains around 70 percent of the variation in NP staffing, consistent with the notion that NP availability varies across EDs and time categories. Conditional on $\mathbf{T}_i$, a wide range of factors, including arriving patients' characteristics, on-duty physicians' characteristics, and a series of other factors that may vary across days (e.g., patient volumes), has virtually no explanatory power for the number of NPs on duty, further supporting quasi-random variation in NP staffing conditional on $\mathbf{T}_i$.

22. We spoke with eight ED physicians and NPs about scheduling; we also distributed a survey to VHA ED administrators about scheduling and received responses from VHA ED administrators in seven states.

23. Per our interviews and surveys with ED providers and administrators, while there may be revisions to on-duty schedules, revisions are infrequent. Further, revisions, if any, are mostly for providers' unexpected personal reasons (e.g., unanticipated family events) that are arguably unrelated to the patients arriving or other conditions at the ED and are typically done in advance of the shifts, making revisions plausibly exogenous.

Conditional Independence.—For our instrument to be valid, the number of NPs on duty must be uncorrelated with patient potential outcomes, conditional on the baseline controls $\mathbf{T}_i$. The institutional features described above support the conditional independence assumption. Further, two sets of empirical evidence provide strong support for this assumption. First, we show that patient observed characteristics are well balanced across the instrument, conditional on $\mathbf{T}_i$. As shown in Figure 2, patient average characteristics are remarkably stable across the instrument, conditioning on $\mathbf{T}_i$. The coefficients are far from being statistically significant and are negligible in magnitude relative to the sample mean. For completeness, Appendix Figure A.4 reports similar balance tests using each of the various patient characteristics included in $\mathbf{X}_i$ as the dependent variable. Even though these characteristics are strong predictors of patient outcomes (F -statistics around 100 for joint significance, even controlling for $\mathbf{T}_i$, see Appendix Figure A.5), there is little significant relationship between our instrument and the broad range of patient characteristics, conditional on $\mathbf{T}_i$.

As a second set of empirical evidence, we examine the stability of our IV estimates under different sets of controls for patient covariates. Specifically, we divide observable patient characteristics into eight groups and estimate separate regressions that control for each of the $2^8 = 256$ different combinations of patient covariates. We show in the empirical results below that controlling for any combination of patient covariates results in virtually no change in our IV estimates, despite these patient covariates being strong predictors of patient outcomes given the substantial sensitivity of the OLS estimates to their inclusion. Following the logic of Altonji, Elder, and Taber (2005), this evidence implies limited selection bias due to either observed or unobserved patient characteristics that predict patient outcomes. In sum, conditional on $\mathbf{T}_i$, there appears to be little relationship between NP availability and patient characteristics.

Exclusion.—While conditional independence supports a causal interpretation of the reduced-form effect, interpreting the IV estimates as identifying the causal effect of being treated by NPs requires an exclusion restriction. That is, the number of NPs on duty impacts patient outcomes only through the patient’s probability of being treated by an NP, not through any other channels. After discussing our main results, we present a series of empirical evidence supporting the validity of the exclusion restriction. We first show that the characteristics of both on-duty physicians and on-duty NPs are well balanced across the number of NPs on duty, conditional on $\mathbf{T}_i$, consistent with quasi-random variation in the instrument as well as our IV estimates being unlikely to be driven by different sets of physicians or NPs available across days. Second, we find no evidence that NP presence influences physician performance. Third, we investigate a series of alternative explanations, finding no evidence indicating a violation of the exclusion restriction.

Monotonicity.—In the presence of heterogeneous treatment effects, we need to assume monotonicity to

interpret IV estimates as a LATE, i.e., the average causal effect among cases induced by the instrument into being treated by NPs. In our setting, monotonicity requires that cases treated by NPs on days with fewer NPs on duty would also be treated by NPs on days with more NPs, and vice versa.

We examine a testable implication of the monotonicity assumption: The instrument and the probability of being treated by NPs should be positively correlated for any subsample defined by patient characteristics. We test this implication in Appendix Figure A.6, where we split the sample by patient characteristics and estimate the first-stage effect separately for each subsample. In particular, we divide the sample by patient age, marital status, gender, race, number of Elixhauser comorbidities, and predicted 30-day mortality.²⁴ Appendix Figure A.6 shows that for all subsamples, the first-stage estimates are positive and statistically different from zero, consistent with the validity of the monotonicity assumption.²⁵

4 Empirical Results

In this section, we present our empirical results. We start by showing the effect of NPs on resource use and downstream patient outcomes. We find that relative to physicians, on average, NPs use more medical resources in the ED visit (longer lengths of stay and higher spending) and exhibit a higher 30-day preventable hospitalization rate. We do not find statistically significant NP effects on 30-day mortality in the overall sample. Next, we examine mechanisms and responses related to the performance difference between NPs and physicians. First, we show that experience may play a role in the NP-physician gap in some outcomes. Second, we find NP responses to lower skill in clinical decision-making, in calling on external resources and setting prescription thresholds. Third, we find that the NP effects on lengths of stay and spending during the ED visit are smaller for less complex and less severe cases. Fourth, we show that patient allocation between NPs and physicians responds toward optimality in the sense of skill-task matching in organizations. Finally, we characterize compliers, present evidence supporting the exclusion restriction, and show a series of additional robustness checks.

24. Predicted 30-day mortality is generated from a linear regression of actual 30-day mortality on the full set of patient characteristics $\mathbf{X}_i$ in Equations (1) and (2).

25. An interesting pattern in Appendix Figure A.6 is that the first-stage estimates appear to be larger for healthier patients: younger patients, those with fewer Elixhauser comorbidities, and those with lower predicted 30-day mortality. This pattern reflects that compliers are more heavily concentrated among healthy patients. In Section 4.4, we compute the characteristics of compliers and never-takers. We find consistent evidence that compliers are healthier than the overall sample, while never-takers are riskier.

4.1 Resource Use

As summary measures of resource use, we start by examining the effect of NPs on patient length of stay and cost of care during the ED visit. Panels A-B of Figure 3 show the reduced-form estimates of patient log length of stay and log cost of the ED visit versus the instrument (i.e., the number of NPs on duty), conditional on the baseline controls $\mathbf{T}_i$. Log length of stay and log cost increase significantly with the instrument. As a comparison, we also plot in the panels patients' predicted log length of stay and predicted log cost, both of which are well balanced across the instrument.²⁶

Columns 1 and 2 of Table 2 report the OLS and IV estimates of the effect of NPs on patient log length of stay and log cost of the ED visit, along with the reduced-form coefficients. All regressions control for the full set of controls described in Section 3.1. The OLS estimates (Panel A) show that cases treated by NPs have significantly shorter lengths of stay and lower costs, which could reflect NPs treating healthier cases than physicians, at least in terms of observable characteristics in Table 1. Exploiting quasi-random variation in patients' probability of being treated by NPs, the IV estimates (Panel C) show that NPs raise patients' medical resource use in the ED visit: On average, cases quasi-randomly assigned to NPs have lengths of stay that are 11 percent longer and ED costs that are 7 percent higher. Given sample means, the NP effects equal an 18-minute increase in the length of stay and a \$66 increase in the cost per ED visit.

Figure 4 examines the robustness of our OLS and IV estimates to the inclusion of different combinations of patient controls. Specifically, we divide patient covariates into eight subsets: (i) five-year age-bin indicators; (ii) marital status; (iii) gender; (iv) race indicators; (v) dummies for 31 Elixhauser comorbidities; (vi) vital signs; (vii) prior health care use; and (viii) indicators for the three-digit primary diagnosis of the visit. We then estimate separate regressions that control for each of the $2^8 = 256$ different combinations of patient covariates for each outcome for both OLS and IV estimations. Figure 4 shows the range of the coefficients across specifications with different patient controls. Each n on the x -axis reports the number of covariate subsets included. For each n , we plot the maximum, mean, and minimum of the estimated coefficients for the effect of NPs using all possible combinations with n (out of eight) subsets of patient covariates.

Figure 4 shows a stark divergence between the OLS and IV estimates. The OLS estimates are negative and decline sizably in magnitude with the addition of patient controls. For example, in the OLS results, conditioning only on baseline controls $\mathbf{T}_i$, we find that cases treated by NPs have 30 percent lower costs than cases treated by physicians. When we add the full set of patient controls, the difference attenuates to 10 percent, a decline of two-thirds. The lower health risks of cases treated by NPs (Table 1) and the sensitivity of the OLS estimates to patient controls suggest the presence of selection bias due to unobservable

26. We form these predictions using linear regressions of actual outcomes on the full set of patient covariates $\mathbf{X}_i$.

patient characteristics. In contrast, the IV estimates are remarkably robust to controlling for any combination of patient covariates: Despite any controls, the IV estimates for the NP effects on length of stay and cost remain stable at 11 and 7 percent, respectively. Following the logic of Altonji, Elder, and Taber (2005), the stability of the IV estimates implies limited scope for selection on either observable or unobservable patient characteristics that predict potential outcomes, further supporting the validity of our instrument.²⁷

In Column 3 of Table 2, we assess NP effects on admissions. In the overall sample, we do not find significant NP effects on admissions, while interestingly, Section 4.3 below shows that NPs significantly increase admissions for severe cases. Appendix Figures A.7 and A.8 show actual and predicted outcomes against the instrument for admissions and the robustness of our estimates for admissions to the inclusion of $2^8 = 256$ different sets of patient controls, respectively.

4.2 Patient Outcomes

We next assess NP effects on downstream patient outcomes, in Columns 4 and 5 of Table 2. The IV estimate suggests a higher 30-day preventable hospitalization rate under NPs: Compared to physicians, NPs, on average, raise the 30-day preventable hospitalization rate by 0.25 percentage points, equivalent to a 20 percent increase relative to the mean. Preventable hospitalization is frequently used as a quality-of-care outcome by both researchers and policymakers (e.g., New Jersey Department of Health 2011; Nyweide et al. 2013; California Health and Human Services Agency 2025; Currie and Zhang 2025). Differences in health risk detection and thus timely preventive treatment and differences in treatment quality conditional on risk detection during the ED visit can lead to differences in post-discharge preventable hospitalization rates. The IV estimate for 30-day mortality is less precise. The 95 percent confidence interval for the estimate shown in Table 2 is -0.34 to 0.11 percentage points.

Panel C of Figure 3 and Panel B of Appendix Figure A.7 plot actual and predicted outcomes against the instrument for patient outcomes in this subsection; Panel C of Figure 4 and Panel B of Appendix Figure A.8 show that the IV estimates for patient outcomes are remarkably robust to the inclusion of $2^8 = 256$ different sets of patient controls, whereas the OLS estimates are highly sensitive to including patient controls.

The results suggest that NPs and physicians, on average, operate on different production functions: NPs, on average, use more inputs in the ED visit (longer lengths of stay and higher costs) but exhibit a higher

27. Another possible explanation for the divergence between the OLS and IV estimates is heterogeneity in treatment effects, since OLS reports the average effect among the analysis sample, while IV reports the average effect among compliers (i.e., cases on the margin of being treated by NPs). To explore this possibility, we follow the procedure by Bhuller et al. (2020) and reweight the analysis sample to match the sample of compliers using predicted 30-day mortality, i.e., a composite index of all patient observables. The OLS estimates with the reweighting still differ in sign from the IV estimates, suggesting that the difference between the OLS and IV estimates cannot be accounted for by heterogeneity in NP effects, at least not by heterogeneous effects across observables.

30-day preventable hospitalization rate. The higher 30-day preventable hospitalization rate also suggests that the greater resource use among NPs is unlikely to be simply because NPs have greater bandwidth to spend time with patients, which we would expect to lead to better patient outcomes under NPs. These results indicate that, when comparing NPs and physicians as two professional classes, physicians, on average, exhibit higher productivity in the ED. However, we note that the estimated average productivity difference between NPs and physicians may be specific to the ED setting, which features, on average, more complex cases than settings such as primary care. Importantly, despite the relative complexity of cases in the ED and the stark differences in training between NPs and physicians, we find substantial overlap between the productivity distributions of NPs and physicians, as shown below.

4.3 Mechanisms and Responses

This section examines mechanisms and responses related to the average performance difference between NPs and physicians. First, we show that experience may play a role in the NP-physician gap in some (though not all) outcomes. Second, we find NP-physician differences in clinical decision-making, in calling on external resources, and in setting prescription thresholds. Third, we show that the NP effect on resource use is smaller for less complex and less severe cases. Finally, we find patient assignment responses toward optimality in the sense of skill-task matching.

Provider Experience.—First, we ask whether experience affects the magnitude of the performance difference between NPs and physicians. The professions of NPs and physicians entail stark differences in the training that new members undergo and in the selectivity of choosing new members. Whether physicians are more productive than NPs due to training or innate ability has important policy implications. While it is difficult to disentangle these two mechanisms, the extent to which the NP-physician performance gap varies with experience may shed suggestive light on this question. If NPs could be made more productive with more extensive training, we may see the gap narrows with experience. On the other hand, if the gap derives from innate ability, we may see the gap persists or even widens with experience.

We form measures of both general and specific experience. We measure general experience as the number of cases the provider has treated from the start of the study period to the day before the current case's visit. We measure specific experience as the number of cases with a condition the same as the current case (measured by the three-digit primary diagnosis) the provider has treated from the start of the study period to the day before the current case's visit. For ease of interpretation, we standardize both general and specific experience to have a mean of zero and a standard deviation of one for NPs and physicians separately.

Our empirical model takes the following form:

$$y_i = \delta_1 \text{NP}_i \times \text{Experience}_i + \delta_2 \text{NP}_i + \delta_3 \text{Experience}_i + \mathbf{T}_i \boldsymbol{\eta} + \mathbf{X}_i \boldsymbol{\beta} + \varepsilon_i. \quad (3)$$

Experience_i denotes standardized experience within each provider type.²⁸ We instrument for NP_i and $\text{NP}_i \times \text{Experience}_i$ using Z_i (i.e., the number of NPs on duty) and $Z_i \times \text{Experience}_i$.

We find that, for some outcomes, experience predicts a smaller NP-physician performance gap. Panel A of Table 3 examines the role of specific experience. For length of stay, Column 1 indicates that a one-standard-deviation increase in specific experience among NPs and physicians is associated with a 5.8 percent decline in the performance gap, reducing the gap at the mean experience levels by about half. The NP-physician cost gap similarly declines with specific experience. Panel B shows that increasing general experience by one standard deviation in both NPs and physicians is associated with a 10 percent decline in the NP-physician gap in length of stay. Unlike length of stay and cost, the NP-physician gap in the 30-day preventable hospitalization rate remains stable across levels of general or specific experience, suggesting that experience may not fully eliminate the NP-physician performance gap.

We examine alternative measures of experience to address measurement concerns. One concern is that, since we may not observe all the cases providers saw from the start of their careers, our measures of experience are imperfect proxies for it. To mitigate this concern, we restrict the sample to cases visiting after 2017, so that our experience measures have at least a one-year look-back window. We also measure experience based on cases seen in the prior year (i.e., 365 days before the index day), so that the estimates precisely represent heterogeneity by prior-year experience.²⁹ As the effect of experience may decay with time (e.g., Benkard 2000), recent experience could be more important than experience gained in the relatively distant past. A second concern is that the number of cases a provider has seen may capture speed, which may affect productivity independent of experience (e.g., faster providers accumulate more cases and discharge patients earlier). To examine this concern, we include an alternative (general) experience measure—the number of days a provider has worked since the start of our study period to the day before the current case’s visit—which is independent of speed. Appendix Tables A.4 to A.6 show that our estimates are qualitatively similar under these alternative measurements.³⁰ Overall, the results suggest that the NP-physician gap in

28. A one-standard-deviation increase in specific experience equals 108 and 107 cases for NPs and physicians, respectively. A one-standard-deviation increase in general experience equals 1,855 and 1,258 cases for NPs and physicians, respectively.

29. Specifically, we measure general experience as the number of cases the provider treated in the preceding 365 days; we measure specific experience as the number of cases with the same three-digit primary diagnosis as the current case the provider treated in the preceding 365 days. We exclude from this estimation cases visiting in the first year of our analysis since we cannot fully observe their providers’ experience in the preceding 365 days.

30. We do not include age as a measure of general experience in this heterogeneity analysis because NPs often practice as nurses

some (though not all) outcomes may narrow with experience, instead of being persistent.

Clinical Decision-Making.—Next, we examine clinical decision-making that may respond to the differences in human capital, specifically, diagnostic skill. Providers with lower diagnostic skill may draw on more external resources, such as consults and diagnostic tests. They may also adjust treatment thresholds for decisions with asymmetric costs between type I and type II errors (Chan, Gentzkow, and Yu 2022). These could be optimal responses to skill differences; they could also increase the cost of care, manifesting in productivity differences.

Columns 1-3 of Table 4 report the effect of NPs on the use of informational resources, using the 2SLS estimation in Equations (1) and (2). Column 1 shows that, relative to physicians, NPs are more likely to use consults: NPs increase consults by 2.6 percentage points, or 11 percent of the sample mean. Columns 2 and 3 show that, relative to physicians, NPs are more likely to order CT scans and X-rays, the two primary diagnostics in the ED setting. NPs increase CT scan and X-ray ordering by 1.2 and 2.0 percentage points, respectively, or 8.3 and 6.9 percent of the respective sample means.

These results suggest that NPs are more likely to collect resource-intensive information from external sources than physicians. This could directly increase lengths of stay and medical costs, since consults and diagnostics take time and resources. On the other hand, consults and diagnostics allow lower-skilled providers to improve decision-making by incorporating information from other experts and, thus, may not be suboptimal. These findings may also imply that NPs seek inputs from other experts when needed, even when allowed to practice independently, which may be reassuring to policy discussions about granting NPs full practice authority.

Next, we evaluate skill from the lens of thresholds for prescriptions with asymmetric costs. Specifically, as shown by Chan, Gentzkow, and Yu (2022), provider skill may correlate with treatment thresholds when the costs of false-positive (type I) and false-negative (type II) errors are asymmetric. Compared to higher-skilled providers, lower-skilled providers may (*optimally*) adjust their treatment thresholds with less information. Specifically, providers with less information may more frequently opt for treatment when false negatives (not treating when a case should have been treated) are costlier than false positives (treating when a case should not have been treated); conversely, lower-skilled providers may less frequently opt for treatment when false positives are costlier than false negatives. Despite more frequent prescriptions, lower-skilled providers may nonetheless incur more false negatives; similarly, they may still incur more false positives with fewer prescriptions.

We choose two important prescriptions with different asymmetries in the costs of type I and type II

for varying lengths of time before becoming NPs and thus, unlike for physicians, age is a noisy measure of experience for NPs.

errors—opioids and antibiotics—and examine prescription thresholds and downstream outcomes. For both of these prescriptions, the clinical indications for appropriate use are often unclear and require clinical judgment.³¹ We estimate the NP-physician difference in prescriptions and downstream outcomes using the 2SLS model in Equations (1) and (2). Column 4 of Table 4 shows that, for opioids, which have higher false-positive costs—e.g., addiction and overdose among patients who should not have received opioids compared to continued pain among patients who should have received them—NPs lower opioid prescriptions by 1.8 percentage points, or 20 percent of the mean. For antibiotics, which have higher false-negative costs—e.g., non-treatment of a potentially life-threatening infection compared to antibiotic resistance—NPs raise antibiotic prescriptions by 4.0 percentage points, 6.3 percent of the mean.³² Yet, as Appendix Table A.7 shows, compared to those treated by physicians, patients treated by NPs are likelier to incur return visits with infections despite higher antibiotic prescribing, and they obtain similar rates of subsequent opioid use disorder despite lower opioid prescribing.

Case Complexity and Severity.—In this analysis, we exploit the wide variety of cases that arrive at the ED to examine heterogeneity in the effects of NPs by case complexity and severity. Following Imbens and Rubin (1997), we estimate complier potential outcomes under NPs and under physicians. Specifically, we estimate complier potential outcomes under NPs using the following IV regression:

$$y_i \cdot \text{NP}_i = \sum_{g=1}^G \mathbf{1}(\text{Group}_i = g) [\delta_g \text{NP}_i + \lambda_g] + \mathbf{T}_i \eta + \mathbf{X}_i \beta + \varepsilon_i, \quad (4)$$

where $y_i \cdot \text{NP}_i$ is the interaction between patient outcome and the indicator for being treated by an NP, $\mathbf{1}(\text{Group}_i = g)$ is an indicator for case i belonging to group $g \in \{1, \dots, G\}$ characterizing complexity or severity. As a natural extension of our main IV model, we instrument for the interactions between $\{\mathbf{1}(\text{Group}_i = g)\}_{g=1}^G$ and NP_i by interacting $\{\mathbf{1}(\text{Group}_i = g)\}_{g=1}^G$ with Z_i , where Z_i is the number of NPs on duty. We estimate complier potential outcomes under physicians using an IV regression similar to Equation (4) but with a dependent variable of $y_i \cdot (\text{NP}_i - 1)$.

We consider two partitions of cases. First, we divide cases into quartiles by their number of Elixhauser comorbidities and refer to higher quartiles as more complex cases. Second, we divide cases into two groups based on whether the condition severity, measured by the 30-day mortality for the three-digit primary diagnosis, is equal to or above the 95th percentile of the sample. The 30-day mortality for this top-severity

31. See, for example, Fleming-Dutra et al. (2016), Huang et al. (2018), Butler et al. (2019), and Neuman, Bateman, and Wunsch (2019).

32. Since opioids apply to a wide range of conditions, we include all patients in examining opioid prescriptions. As antibiotics generally only apply to patients with infections, we restrict the sample to patients with respiratory or genitourinary system infections, i.e., two common types of infections.

group is 8.6 percentage points, compared with 1.2 percentage points for the whole sample. Included in this group are cases with severe conditions, such as acute kidney failure and acute myocardial infarction (AMI), which are potentially more challenging to manage.³³

We find the NP effects on lengths of stay and ED costs are smaller when cases are less complex and less severe, as shown in Figure 5. For example, for cases in the second complexity quartile, NPs increase their ED costs by 7 percent; for cases in the highest-complexity quartile, NPs show a 12 percent increase in ED costs; and for cases with a condition severity at or above the 95th percentile, NPs show a 25 percent increase in ED costs. Table 5 summarizes the heterogeneous NP effects.³⁴ The table further examines four severe conditions: sepsis, stroke, AMI, and heart failure. We find smaller NP effects on length of stay and ED costs among cases without these severe conditions. Appendix Table A.9 reports the outcome mean for each group of cases in Table 5.

While NPs do not affect admissions for low-risk patients, they raise admissions among high-risk patients. For cases with severity at or above the 95th percentile, NPs show an increased admission rate of 26 percentage points, or 42 percent compared to the mean. For sepsis, stroke, AMI, and heart failure cases, NPs increase their admission rates by 45, 73, 123, and 20 percentage points, respectively, or 50, 130, 214, and 31 percent relative to the mean. From the perspective of allocative efficiency, the greater use of admissions and other resources by NPs relative to physicians for more severe cases may not be suboptimal: It may improve NP performance on patient health outcomes relative to a counterfactual scenario that uniformly constrains NPs to the same admission rates as those of physicians.

We find interesting evidence of this tradeoff between resource use and patient health outcomes when we examine heterogeneity in NP effects on 30-day preventable hospitalizations. In particular, we find decreasing NP effects on 30-day preventable hospitalization with case complexity or severity (Column 5 of Table 5). If NPs are less skilled at treating more complex or more severe cases, they may worsen outcomes more for these cases. On the other hand, as NPs increase their intensity of care by a larger extent for these cases, the higher incremental care may in fact reduce preventable hospitalizations. As shown in Panel A of Table 5, for cases in the highest-complexity quartile, NPs increase lengths of stay and ED costs by 28 percent and 12

33. Appendix Table A.8 summarizes the 10 most common three-digit primary diagnoses in this group. The largest diagnosis category is heart failure, followed by acute kidney failure. The top 10 diagnoses also include AMI, a form of sepsis, and a form of respiratory failure. The 30-day mortality rates range between 5 and 17 percentage points. A related question is why severe cases are assigned to NPs. While the data show a reduced probability of NP assignment among riskier patients, the probability remains positive (the share of patients assigned to NPs among patients whose condition severity is below and at least as high as the 95th percentile of the sample is, respectively, 0.24 and 0.11). Possible reasons that severe cases are assigned to NPs may include (i) triage providers may misevaluate patient severity (Chan and Gruber 2020) and assign severe cases to NPs, and (ii) when physicians are occupied with earlier cases, it may be more efficient to assign severe cases to available NPs than to delay care.

34. Results in this table are estimated using Equation (4) with patient outcome y_i as the dependent variable. By construction, this dependent variable is the difference between the dependent variables used to estimate potential outcomes: $y_i = y_i \cdot NP_i - y_i \cdot (NP_i - 1)$.

percent, respectively, without a significant effect on the 30-day preventable hospitalization rate. For cases in the lower-complexity quartiles, NPs reveal a smaller-magnitude, and sometimes insignificant, effect on length of stay and medical costs, yet NPs significantly raise the 30-day preventable hospitalization rate for these cases. For the most severe cases (i.e., those with severity at or above the 95th percentile), NPs *reduce* 30-day preventable hospitalization by three percentage points, while increasing length of stay by 99 percent, ED costs by 25 percent, and admissions by 26 percentage points (Panel B of Table 5).

Patient Assignment.—Finally, we examine patient assignment between NPs and physicians. The evidence on case heterogeneity above and, more directly, the counterfactual cost estimates by case complexity in Section 5 below, suggest a comparative disadvantage for NPs in treating complex and severe cases. We do not find a set of cases in which NPs outperform physicians, suggesting that NPs are also at an absolute disadvantage in the ED setting. These stylized facts imply a qualitatively optimal assignment of patients to NP versus physician provider classes, in the sense of skill-task matching in organizations (Acemoglu and Autor 2011): Of patients available, NPs should receive the healthier ones, and NPs should receive fewer patients when physicians are more available. Some recent studies suggest that optimal worker-job matching may not hold in practice (e.g., Adhvaryu, Kala, and Nyshadham 2022; Minni 2023). Yet empirical patient allocation between NPs and physicians consistent with the direction of optimal assignments can augment the efficiency of using NPs.

Appendix Figure A.9 provides descriptive insights into patient assignment by exploiting variation in NP staffing and patient arrivals, conditional on our baseline controls. Panels A-C show that NPs overall are assigned healthier cases, consistent with Table 1. The average complexity and severity of cases assigned to NPs, as measured by patient age, comorbidities, and predicted 30-day mortality, increase with NP staffing, while the average complexity and severity of all cases remain stable. The pattern suggests an assignment process in which the first cases assigned to NPs have the lowest health risks; when more NPs (and fewer physicians) are available, cases incrementally assigned to NPs are riskier compared to those initially assigned to NPs but are still relatively healthy among the remaining cases to be assigned.³⁵ Panel D of the figure assesses the probability that a patient is assigned to an NP as a function of ED busyness measured by the number of other patients arriving in the analysis time window (i.e., 8 a.m. to 6 p.m.) of the ED-day cell. We find a small but clear trend in which patients are more likely to be assigned to NPs when the ED is busier, consistent with the efficiency of assigning fewer patients to NPs when physicians are less occupied.³⁶

35. As a result, cases left for physicians become riskier with more NPs. A potential question is whether the increased risk of cases assigned to physicians on days with more NPs affects physicians' performance, violating an exclusion restriction. To address this concern, Appendix Table A.10 controls for multiple measures of average health risk (age, comorbidities, and predicted 30-day mortality) for cases assigned to physicians within the ED-day cell. The results show that our IV estimates are highly robust.

36. This relationship is qualitatively the same when further conditioning on the number of NPs and the number of physicians

Taken together, the empirical patient allocation between NPs and physicians is consistent with the direction of optimal skill-task matching in organizations, which may augment the efficiency of using NPs.

4.4 Complier Characteristics

Our IV estimates represent the LATE, i.e., the average causal effect among complier cases quasi-randomly assigned to NPs versus physicians due to the instrument. To better understand this LATE, we examine complier characteristics following the approach developed by Abadie (2003), as described in Appendix A.2. Appendix Table A.11 reports the results. Consistent with NPs treating less severe cases, compliers are healthier than the average case. Compared with the average case, compliers are younger and have fewer Elixhauser comorbidities, fewer inpatient stays in the prior year, and lower predicted mortality. Appendix Table A.11 also examines characteristics of never-takers (of NPs), following an approach from Dahl, Kostøl, and Mogstad (2014) that we detail in Appendix A.2. In line with the notion that NPs treat healthier cases than physicians, never-takers are riskier than the average case, and both are riskier than compliers. (Our setting contains no always-takers since patients cannot be assigned to NPs on days with no NPs.)

4.5 Exclusion Restriction

As discussed in Section 3.2, interpreting the IV estimates as the causal effect of being treated by NPs requires the number of NPs on duty to affect patient outcomes only through patients' probability of being treated by NPs, not through any other channels. In this section, we present evidence supporting the exclusion restriction. We first show that the characteristics of both on-duty physicians and on-duty NPs are well balanced across the number of NPs on duty. Second, we assess and find little support for the possibility of productivity spillovers that NP presence influences physician performance in our setting. Third, we examine a range of factors that may vary across days, including patient volume, total staffing, and patient wait times, and consider patient-physician gender concordance, finding no evidence to suggest these factors are driving our IV estimates.

Balance in Provider Characteristics.—To start, we investigate whether on-duty physicians are similar across days with different numbers of NPs. If such a balance does not hold, our IV estimates could be driven by compositional changes of physicians. Figure 6 reports the balance for various physician characteristics. Specifically, we consider an ED-day level analysis that asks whether on-duty physicians' staffing the ED on that day. A related question is whether the increasing NP assignment with ED busyness may affect the estimated NP effects. In contrast to OLS, a correlation between NP assignment and busyness per se would not affect estimates from our IV approach. In Section 4.5, we also show that our IV estimates are highly robust to controlling for the number of patients arriving.

average characteristics (weighted by the number of cases treated by each physician) are independent of the number of NPs on duty, conditional on ED-by-year, ED-by-month, and ED-by-day-of-the-week indicators.³⁷

We examine three sets of physician characteristics: (i) demographics of age and gender; (ii) measures of physician “value-added,” reflecting physician risk-adjusted impact on patient 30-day mortality; and (iii) measures of physician “practice style,” reflecting a physician’s risk-adjusted average input choices in terms of length of stay and ED costs (see Appendix A.3 for construction details). Figure 6 shows that each of these physician characteristics is balanced across the instrument, conditional on the baseline controls.

We similarly examine whether NP characteristics are systematically different across days with differing numbers of NPs on duty. If such systematic variation exists, our baseline IV strategy may not be able to disentangle the effect of being treated by NPs from that due to the potentially different quality of NPs across days.³⁸ Following Figure 6, Figure 7 shows that an analogous set of NP characteristics is well balanced across days with different numbers of NPs, conditional on baseline controls.³⁹

The balance in on-duty physicians’ and on-duty NPs’ characteristics across the instrument also supports quasi-random variation in the instrument, as with the balance in patient characteristics. Otherwise (e.g., if the variation is driven by patient or staffing shocks), we would observe systematic changes in these characteristics with NP staffing.

Assessing Productivity Spillovers.—We then consider the possibility of spillovers between NPs and physicians: Whether NP presence may influence physician performance. For example, if NPs ask physicians in the ED for assistance, they may slow down physicians. Alternatively, a change in peers from days without any NPs to days with NPs may influence physician performance as, e.g., physicians may come under different degrees of peer pressure that motivate them to work differently (Chan 2016; Silver 2021).

However, we find little empirical evidence to suggest meaningful spillovers between NPs and physicians. First, if NPs ask physicians for assistance, we may expect the outcomes of patients treated by NPs to depend on the quality of physicians on duty. Using the value-added measure described in Appendix A.3, we find no such

37. The empirical specification takes the form $\bar{y}_{jd} = \tilde{\lambda}Z_{jd} + \tilde{\mathbf{T}}_{jd}\tilde{\eta} + \varepsilon_{jd}$, where $\bar{y}_{jd}$ is the average characteristics of physicians on duty at ED j on day d (weighted by the number of cases treated by each physician), Z_{jd} is the number of NPs on duty, and $\tilde{\mathbf{T}}_{jd}$ includes ED-by-year, ED-by-month, and ED-by-day-of-the-week indicators. We cluster standard errors by ED. Since our main sample has only 44 EDs, a potential question is whether the estimated standard errors are biased, given the relatively small number of clusters. While there is currently no clear-cut definition of “small,” if anything, such a potential issue would bias us toward rejecting the null hypothesis of physician balance across the instrument (Bertrand, Duflo, and Mullainathan 2004; Cameron and Miller 2015). Nonetheless, as a robustness check, we apply the correction for the small number of clusters using the wild cluster bootstrap as suggested by Cameron and Miller (2015); we find no meaningful change in the standard error estimates.

38. This concern applies to our baseline IV strategy since it uses variation in the number of NPs in addition to whether there is an NP on duty. In Section 4.6, we include alternative estimations using variation in the extensive margin of whether there is any NP on duty and restricting the sample to days with zero or only one NP on duty; results are virtually unchanged.

39. Though the balance-test coefficient for physician practice style in terms of ED costs has a p -value below 0.1 at 0.097, given multiple hypothesis testing, it can reflect a chance finding. To further assess the robustness, Appendix Table A.12 reports our IV estimates controlling for on-duty providers’ practice styles, showing that the results are stable.

relationship in Appendix Table A.13. Second, with spillovers—either through assistance or peer pressure, or importantly, any other channels by which NP staffing may affect physician performance—outcomes for patients treated by physicians could change with the presence of NPs. Directly regressing outcomes of patients treated by physicians on the NP presence suffers from patient selection since physicians are allocated riskier cases on days with NPs (as healthier cases are assigned to NPs). To circumvent this issue, we look at cases arriving between 5 and 8 a.m., i.e., patients who arrive before the typical start of NP shifts so that they are unlikely to be assigned to NPs, but whose stay overlaps with NP shifts so that their physicians could be subject to spillover effects from NPs.⁴⁰ As shown in Panel A of Appendix Table A.14, we find no evidence of spillovers from NP presence on physicians’ patients overall. In Panel B, we examine days with high workloads, when NP spillovers may be more detectable, and again find no evidence of spillovers in this subsample.⁴¹ In sum, there appears to be little evidence that NP presence affects physician performance in our setting.

Robustness to Additional Factors.—Finally, we show the robustness of our IV estimates to factors that may vary across days, including total physician equivalents on duty, the total number of cases arriving, and patient wait times (the time between arrival at the ED and assignment to a treating provider). We also show that our IV estimates are unlikely to be driven by patient-provider gender mismatch.

We control for total physician equivalents on duty to mitigate the concern that the effective level of staffing may vary across days with different numbers of NPs on duty. Conditional on T_i , the ED plausibly maintains a consistent level of total effective staffing across days, lending credence to our baseline IV estimates. The coefficient in Panel A of Figure 1 being less than one in magnitude does not necessarily indicate increased staffing with more NPs, since the number of providers is not a sufficient statistic for staffing. For example, given that NPs take longer to discharge a patient but do not handle more or sicker patients simultaneously, the NP-physician substitution rate is below one, and maintaining balanced total effective staffing across days with differing NP availability would require a coefficient less than one in magnitude in Panel A of Figure 1. Importantly, if an increased number of NPs is associated with higher total staffing, we would expect better patient outcomes, yet we find a higher 30-day preventable hospitalization rate on days with more NPs. For further robustness checks, we control for physician equivalents on duty, calculated as the sum of the number

40. To restrict further the possibility of patient selection between NPs and physicians, we exclude patients arriving between 5 and 8 a.m. in ED-day cells with any patients assigned to NPs. As shown in Columns 1 and 2 of Appendix Table A.14, patient health risks are balanced between days with and without NPs in this analysis sample.

41. A potential question is whether the increased average health risk of patients assigned to physicians on days with NPs affects physicians’ overall performance, violating an exclusion restriction. To assess this concern, as described in footnote 35, Appendix Table A.10 controls for multiple measures of the average health risk (age, number of Elixhauser comorbidities, and predicted 30-day mortality) of patients assigned to physicians in the ED-day cell. The results show that our IV estimates are remarkably stable, suggesting that the different sets of patients assigned to physicians on days with NPs are unlikely to affect our IV estimates.

of on-duty physicians and the number of on-duty NPs multiplied by the substitution rate between NPs and physicians. Since we do not know how many physician equivalents an average NP represents, we assume a series of substitution rates—0.25, 0.5, 0.75, and 1, assuming that one NP substitutes for 0.25, 0.5, 0.75, and 1 physician, respectively. We cap the substitution rate at one since, as described above, one NP appears insufficient to substitute for one physician. When the assumed substitution rate is higher than the actual rate, the IV estimates will be biased by lower staffing, leading to underestimated IV estimates for log length of stay and ED cost when lower staffing rushes providers to shorten patient visits and reduce treatments (Silver (2021) shows that when ED physicians are rushed, they discharge patients early and reduce treatments). Conversely, if the assumed substitution rate is lower than the actual rate, the IV estimates will be confounded by higher staffing, potentially leading to overestimated NP effects on log length of stay and ED cost. The more the assumed substitution rate is above (below) the actual rate, the smaller (larger) the estimates will be relative to the truth. Similarly, if lower (higher) staffing worsens (improves) patient outcomes, e.g., mortality, we would see decreased IV estimates for mortality with lower assumed substitution rates. Since the interval [0.25, 1] plausibly covers the true substitution rate, we consider this robustness test a bounding approach. Appendix Figure A.10 reports the results. Though the IV estimates modestly change with the substitution rate as explained above, the estimates are qualitatively similar to our baseline IV estimates.

Turning to the total number of cases arriving and wait times, while potentially important, Appendix Table A.15 shows that these two factors do not affect our estimates: the IV estimates are highly robust to controlling for these factors. Since patient wait times are potentially endogenous (healthier cases could be assigned a lower priority and thus wait longer), we instrument for wait times using the average wait time of cases visiting on the same day at the same ED as the index case in Appendix Table A.15.

We also ask if the estimated NP effect is driven by patient-provider gender mismatch since the vast majority of patients are male (91 percent), physicians are primarily male (74 percent), and NPs are mostly female (79 percent). Appendix Table A.16 explores this possibility by asking whether the effect of NPs varies by whether the patient's and provider's genders match. The results show little heterogeneity.

4.6 Additional Robustness Checks

Appendix Tables A.17-A.21 report additional robustness checks. Appendix Table A.17 shows that our findings are stable with alternative standard error clustering approaches: clustering by ED-day cell or two-way clustering by ED-day cell and provider. Panels B and C of Appendix Table A.18 show the robustness of our estimates to, respectively, an alternative count of on-duty NPs that includes any NP with at least one case (instead of at least two cases as in our baseline IV) in the analysis time window of an ED-day cell and

another count that leaves out the index case in measuring on-duty NPs. In Panels D-F of Appendix Table A.18, we construct three alternative instruments: the share of on-duty providers who are NPs, an indicator for any NPs on duty, and the share of cases assigned to NPs (leaving out the index case) in the ED-day cell; we find the results are stable. Appendix Table A.19 shows that our results remain similar when looking at the margin between days with no versus only one NP on duty. Appendix Table A.20 shows that our results for hospital admissions in the ED visit and 30-day preventable hospitalizations are robust to considering hospital stays outside of the VHA.⁴² Appendix Table A.21 expands the sample to all 110 VHA EDs that use NPs—which doubles the number of cases and triples the number of providers in our sample—and shows that the results are similar.⁴³

5 Counterfactual Scenarios

In this section, we consider two counterfactual policy scenarios that quantify the average NP-physician performance difference in financial terms. First, we consider the cost implications of substituting physicians with NPs. Second, we consider augmenting the existing supply of physicians with NPs.

We first perform a simple calculation of the cost of assigning 25 percent of all cases in VHA EDs to NPs instead of physicians, approximately the share of cases treated by NPs in our sample, which consists of EDs that are early adopters of NPs. Using the 2SLS specification in Section 3.1, we estimate the effect of NPs on log total spending per case, where total spending is the sum of the three main cost components investigated in Section 4: the cost of ED care, hospital admission, and 30-day preventable hospitalization.⁴⁴ Assuming that the effects of NPs across all EDs are similar to those in our sample, we multiply the extra spending per case associated with NPs by the total number of cases in the 25 percent set. Based on point estimates, we find an extra cost of \$197 million per year for the VHA.⁴⁵ This figure is 2.6 times the yearly NP wage costs

42. Since a share of patients has health insurance coverage in addition to the VHA's (mainly Medicare), we report robustness checks that include hospital stays outside of the VHA for patients who enroll in both the VHA and traditional Medicare. Note that this is likely an upper-bound estimate of the possible bias for our sample since patients without non-VHA health insurance are much less likely to have hospital stays outside of the VHA. Regardless, the estimates are stable.

43. As described in Section 2.3, to focus on a single margin between NPs and physicians, our main analysis focuses on EDs that use only NPs and physicians. In Appendix Table A.21, we include all VHA EDs that use NPs regardless of using physician assistants (but exclude cases in ED-day cells with physician assistants).

44. To calculate the cost of hospital admissions and 30-day preventable hospitalizations, we apply the cost estimate of \$19,220 per VHA hospital stay, based on the average length of stay per hospitalization at the VHA and costs per VHA inpatient day reported by the VHA's Health Economics Resource Center (2021).

45. As this estimation focuses on preventable hospitalization effects within 30 days of the ED visit, the extra cost estimated can be viewed as that within 30 days of the ED visit. The NP impact on total spending per case may be larger when we consider costs further downstream. The 95 percent confidence interval of the extra cost in this counterfactual is [-\$16 million, \$411 million]. The *p*-value of the estimate is 0.067. The relatively wide confidence interval mainly stems from hospital admissions (as shown in Panel C of Table 2, the NP effects on ED costs and 30-day preventable hospitalizations are much more precise). In this counterfactual and Section 6 below, we take the societal perspective and view extra medical spending as costs. While this may reflect the perspective of a vertically integrated delivery system such as the VHA, other hospitals in the US may not fully internalize the societal cost of

that the VHA would encounter to assign 25 percent of its ED cases to NPs.⁴⁶

The calculation ignores potential changes in provider wage costs, which may be large, as the average NP wage is only half the average physician wage (Bureau of Labor Statistics 2021a; 2021b). If two NPs substitute for one physician, there may be no wage savings when substituting physicians with NPs. For a conservative estimate, we consider the scenario where one NP may substitute for one physician. Under this scenario, we also arrive at a net cost, \$129 million per year, for assigning 25 percent of VHA ED cases to NPs.⁴⁷

Using LATEs, this analysis assesses counterfactual outcomes for compliers, likely to encompass the 25 percent of cases assigned to NPs. Yet we also consider a counterfactual scenario where EDs assign the least complex 25 percent of cases (by the number of Elixhauser comorbidities) to NPs. Applying the estimates for the lowest-complexity quartile cases in Table 5, we find a lower extra medical cost of \$97 million per year to the VHA. This cost can be reduced to \$29 million per year by lower NP wages when an NP can substitute for one physician, only about one-fifth of the net cost (\$129 million per year) in the counterfactual above.

Despite possible extra spending, our findings do not imply that NPs are inefficient to use. When physician capacity is limited, hiring NPs to augment provider supply and reduce wait times may, nevertheless, be necessary and efficiency-improving. To examine this concept, in Appendix A.4 we consider the trade-off between reducing wait times and increasing resource use with additional NPs. As overcrowding is a significant issue in the ED, wait time has been an important object of attention for policymakers and ED management alike (e.g., Institute of Medicine 2006; American College of Emergency Physicians 2016). We find that roughly nine-tenths of the additional spending to reduce wait time by hiring NPs comes from the increased medical resource use associated with NPs; only one-tenth comes from additional NP wage costs. However, the extra medical cost decreases with case complexity: It declines by about one-third in the lowest-complexity quartile of cases compared to the average case.

Finally, beyond the scope of this paper, there may be important long-run efficiency gains from increased NP use. For example, broader NP adoption may impact physician wages, enable greater physician specialization in complex cases or better coordination between NPs and physicians, or influence the composition of individuals entering the NP and physician professions. These may be fruitful areas for future research.

hiring low-productivity workers since they may not bear the cost of downstream outcomes. They may even receive profits from extra care utilization, encouraging a wage-productivity relationship different from the societal perspective.

46. For this estimation, we divide the total number of cases in the 25 percent set by the average caseload of NPs in our sample, finding that 654 NPs would be needed for treating 25 percent of VHA's ED cases annually. We then multiply the number of NPs needed by the NP wage reported by Bureau of Labor Statistics (2021a), yielding a total wage estimate of \$74.9 million per year.

47. We use the physician and NP wages reported in Bureau of Labor Statistics (2021a; 2021b) for nationally representative wage estimates. Since the Bureau of Labor Statistics does not provide separate wage estimates for NPs in the ED, we apply the wages of the average physician and the average NP. Using instead the wage estimates in the VHA data for NPs and physicians in the ED, we arrive at similar net costs of \$109 million per year.

6 Productivity within Professions

Up to this point, we have focused on estimating the average productivity difference between the professional classes of NPs and physicians. In this section, motivated by growing evidence of productivity variation across providers (within professions), we measure the distribution of productivity within each profession and ask how this intra-professional variation in productivity compares with the difference in productivity between professions, specifically NPs versus physicians.⁴⁸ We find productivity variation within each profession several-fold larger than the productivity difference between the two professions. This leads to substantial overlap between the productivity distributions of NPs and physicians: A large share of NPs perform even better than the average physician. Using detailed administrative data from the VHA, we also evaluate the extent to which productivity within each profession relates to the average health risk of cases assigned and the wages paid to individual providers, finding limited relationships.

6.1 Distribution of Productivity

We first examine the distribution of productivity within each profession. We operationalize this examination by focusing on a measure of the total cost per case. Specifically, for each case, we aggregate the three main cost components investigated in Section 4, the same components we considered in Section 5: the cost of ED care, hospital admission, and 30-day preventable hospitalization.⁴⁹ We then estimate provider effects on the log of this measure of the total cost of each case, with higher effects indicating *lower* productivity. In practice, hospitals may not fully internalize the societal cost of using higher-spending workers, and extra care utilization may be an opportunity to reap additional revenues, particularly for hospitals outside of the VHA. In this paper, we take the societal perspective: We view extra spending as costs and define higher spending effects as lower productivity.

To account for the provider-specific selection of patients in estimating provider spending effects, we use a just-identified IV model that instruments for indicators for treating providers with indicators for providers on duty in the ED-day cell of the patient’s visit, leveraging the quasi-random variation in on-duty providers (conditional on $\mathbf{T}_i$). The model also controls for the baseline control vector $\mathbf{T}_i$ and for robustness checks controls for patient characteristics $\mathbf{X}_i$, as our baseline IV model in Equations (1) and (2). Appendix A.5.1 provides details of the empirical specification and shows that the instruments (i.e., the indicators for on-

48. Doyle, Ewer, and Wagner (2010) show differences in resource utilization decisions among physician trainees, potentially driven by human capital. Abaluck et al. (2016) document variation in physician testing thresholds in the context of pulmonary embolism, and Silver (2021) examines returns to time spent on patients by ED physicians and variation in the physicians’ productivity. Currie and MacLeod (2017) and Currie and MacLeod (2020) show physician skill variation in the contexts of C-section and depression treatments, respectively. Chan, Gentzkow, and Yu (2022) demonstrate important variation in diagnostic skill among radiologists.

49. We multiply the latter two components by the average cost of a hospital stay, \$19,220.

duty providers) are strongly predictive of the treating providers but are independent of arriving patients' characteristics conditional on $\mathbf{T}_i$, supporting the validity of these instruments.

Appendix Table A.22 reports estimates of the variance of provider effects on the log total cost of the ED visit defined above. Using a split-sample approach to account for finite-sample estimation error, we find a variance of 0.045 for physicians and 0.048 for NPs (see Appendix A.5.2 for details of the split-sample approach). These estimates suggest large variation in provider effects: A one-standard-deviation costlier physician and NP increase the total cost of the ED visit by 21 and 22 percent per case, respectively, which are about three times the average NP effect of 6.7 percent on the total cost per case defined above from the 2SLS model in Equations (1) and (2). That is, the productivity variation within each professional class is much larger than the difference in productivity between the two classes. Related to the wide variation in provider effects within each professional class, Appendix Figure A.11 shows that the NP effect varies considerably across EDs (details in Appendix A.6). Accounting for sampling error, the standard deviation of the ED-specific NP effect for each outcome is similar in magnitude to the average NP effect reported in Section 4.⁵⁰

We next investigate the full distributions of provider effects on the log total cost per case, applying a non-parametric empirical Bayes deconvolution approach adapted by Kline, Rose, and Walters (2022) from Efron (2016). This approach extracts a flexible empirical Bayes prior distribution of population provider effects, using the estimated provider effects and their standard errors from the just-identified IV model described above. Specifically, the procedure first yields a distribution of provider z -scores (the ratio of each provider effect to its standard error); it then estimates the distribution of provider effects from the z -score density function. Appendix A.5.3 describes details of the estimation. We apply this procedure separately for NPs and physicians and ensure that the difference between the means of the deconvolved distributions for NPs and physicians equals the NP effect from the 2SLS model in Equations (1) and (2).⁵¹

Panel A of Figure 8 displays the deconvolved density of provider effects. For interpretation, we convert provider effects on log total spending per case on the x -axis to annual spending (not including provider wages) by incorporating the empirical distribution of log total spending per case and the average number of cases a provider treats per year (see Appendix A.5.3 for details of the conversion). Within each professional

50. Consistent with the results in Section 4, we find that provider productivity is correlated with professional class. Within professions, we find no significant correlation between provider productivity and provider observable characteristics (age, gender, whether born in the US, whether receiving medical training in the US (observable for physicians only), and whether having a doctoral degree (for NPs only)). This pattern is consistent with the education literature, which suggests that teacher value-added is not strongly predicted by observables (e.g., Rivkin, Hanushek, and Kain 2005; Jackson, Rockoff, and Staiger 2014). Yet we do not exclude the possibility that provider productivity may be more predictable with a richer set of characteristics.

51. The average difference in total spending per case based on the 2SLS model and the difference between the means of the NP and physician effects from the just-identified IV model weighted by the number of cases seen by each provider are close: 6.7 versus 7.2 percent.

class, the productivity distribution implies greater non-wage spending of about \$910,000 and \$870,000 per year under a provider at the 25th percentile of productivity than under a provider at the 75th percentile for NPs and physicians, respectively—both of which are as high as about three times the mean annual non-wage spending difference between NPs and physicians.

Finally, using the deconvolved density of NP and physician effects on log total spending per case, we estimate the probability that a randomly drawn NP is costlier than a randomly drawn physician (Appendix A.5.3 provides details of the estimation). We find that the probability that a randomly drawn NP is costlier than a randomly drawn physician, as measured by the effect on the total spending for the ED visit defined above, is only 62 percent. Put differently, the probability that a randomly drawn NP is less costly than a randomly drawn physician is as high as 38 percent. This statistic remains large when we adjust the deconvolved productivity distributions to account for possible differences in treatment effects between the overall population and compliers: When we assume the average treatment effect is as large as twice the LATE estimate (higher than the LATE among the highest-complexity quartile patients), the probability that NPs are less costly remains large at 28 percent.

Taken together, the findings in this section suggest a nuanced role for professions in determining workers' productivity. Despite stark differences in training and the relative complexity of cases in the ED setting, there is still substantial overlap in the productivity distributions of NPs and physicians: A large share of NPs may perform even better than the average physician. Therefore, despite the average productivity difference between the two professions, our findings do not necessarily imply inefficiency in using NPs: Not only could NPs provide a valuable labor supply for health care, but the large productivity overlap between NPs and physicians also suggests scope for efficiency gains beyond organizing work solely around professional lines. In addition, the large productivity overlap between NPs and physicians, even in the ED setting, may shed a positive light on the use of NPs in less complex settings (e.g., primary care).

6.2 Relationship with Case Assignment and Wages

Next, we evaluate the distribution of assigned patient risks and paid wages across providers within each professional class. In Panel B of Figure 8, we show the average patient risks, measured by predicted 30-day mortality, across providers in each professional class. While there is considerable overlap in the distributions for NPs and physicians, the difference in means between the two distributions is clear. In Panel C of Figure 8, we plot the distributions of annual wages across providers.⁵² The panel shows substantial wage variation

52. For each provider, we access detailed payment records of the full-time equivalents and wages for each pay period from 2011 to 2020. We convert these data to annualized provider wages by (i) inflation-adjusting payments in any year to corresponding payments in 2020 dollars, (ii) computing a per-hour wage by dividing the sum of (inflation-adjusted) payments by the sum of work hours

within each profession: The gaps between the 10th and 90th percentiles of annual wages are about \$48,000 and \$117,000 for NPs and physicians, respectively. Despite the large wage variation within each profession, there is virtually no overlap in the wage distributions between the two professions.

In Appendix Figure A.12, we present the receiver operating characteristic (ROC) curves of provider productivity, case assignments, and wage payments as characteristics that may distinguish NPs and physicians (details in Appendix A.5.4). The area under the curve (AUC) values are 0.62, 0.75, and 0.99 for productivity, case assignments, and wage payments, respectively. Put differently, assigned patient risks are more predictive of professional class than productivity is; yearly wages are extremely predictive of professional class. Productivity, which exhibits the largest overlap between the distributions of NPs and physicians (Figure 8), has the poorest performance in classifying providers.

Finally, given the potential for large efficiency gains under policies using provider-specific productivity, we explore how our measure of provider productivity (i.e., the reverse of provider effects on the log total spending per case) relates to assigned patient risks and wages paid to each provider. To account for sampling error, we compute the empirical Bayes posterior mean for each provider's effect on log total spending per case (details in Appendix A.5.5). In contrast to the clear gradient between the professional classes of NPs and physicians, we find that provider productivity within each professional class has little bearing on the average predicted 30-day mortality of cases assigned to each provider (Panel A of Appendix Figure A.13).⁵³ Similarly, we find that, within each professional class, higher productivity is not correlated with higher wages (Panel B of Appendix Figure A.13).⁵⁴ Taken together, the evidence indicates that organizations make little use of provider productivity within each professional class when assigning cases or setting wages for providers.

7 Conclusion

Professionals perform some of the most important tasks across a variety of economic sectors. In turn, professional groups play a central role in determining the division of professional labor, the selection and training of future members, and the economic returns to professional work. However, very little is known

across all pay periods, where each pay period covers two weeks, and (iii) multiplying this figure by 26 pay periods and 80 hours per pay period (the number of hours in a full-time-equivalent pay period).

53. In an analysis that divides cases into low- versus high-risk by whether predicted 30-day mortality is below versus above the sample median and separately identifies provider effects on low- and high-risk patients, we find that the gap in spending effects between low- and high-productivity providers is greater for riskier patients. This implies that assigning riskier patients to higher-productivity providers could be efficient.

54. For NPs, though Panel B of Appendix Figure A.13 shows a relationship between wages and productivity, the relationship is only marginally significant and is opposite-signed: NPs with higher effects on spending, i.e., NPs who are lower-productivity, are paid higher wages.

empirically about productivity differences between distinct professions performing overlapping tasks and how they compare to within-profession productivity variation, largely because professions exclude other groups from providing tasks within their “jurisdictions” (Abbott 2014).

In this paper, we exploit a unique opportunity to study two starkly different professional groups—nurse practitioners (NPs) and physicians. Our empirical setting allows us to study the quasi-experimental assignment of ED patients to NPs versus physicians within the Veterans Health Administration (VHA), in the context of a directive granting NPs full practice authority. We use the quasi-random arrival of patients at the ED between times that may differ in the availability of NPs on shift, which drives the probability of being treated by an NP versus a physician. In the ED setting, compared to physicians, on average, NPs use more resources during the visit (longer lengths of stay and higher spending) and exhibit a higher 30-day preventable hospitalization rate.

We further examine behavioral mechanisms and responses to productivity in our setting. We show that experience may play a role in the NP-physician performance gap. We find clinical decision-making in response to skill differences: NPs are likelier to call on external information, and they shift their prescribing decisions to avoid costly errors. We also find that the performance gap between NPs and physicians lessens when cases are less complex and less severe. In response, NPs receive healthier cases and take on a smaller share of the caseload when the ED is less constrained to meet demand. While we find in the ED that, on average, the resource cost implied by the lower productivity of NPs versus physicians outweighs salary savings from hiring NPs instead of physicians, this cost declines when NPs are assigned healthier patients. These patterns suggest returns to training, decision support, and optimal case assignment for NPs.

We find striking productivity variation within each profession, several-fold larger than the difference between the two professions. This leads to substantial overlap between the productivity distributions of NPs and physicians, such that a large share of NPs outperform the average physician. These findings imply that the professional class is, at best, a coarse indicator of productivity. While NPs and physicians receive qualitatively different patient risks and starkly different wages, within each profession, tasks and wages bear limited relationships with worker productivity. The large productivity overlap between the two professions suggests that organizing work solely around professional lines may forego large gains in efficiency.

Considered together, our findings paint a nuanced picture of the role of professions in determining and revealing workers’ productivity, tasks, and wages. Intensive professional selection and training may imply average productivity differences between professional classes that justify sizable wage differences in industries such as health care, where worker judgments may influence the utilization of considerable resources. Nevertheless, professional institutions are likely more effective at separating wages and organizing

tasks between professions than at standardizing productivity within professions. These relationships may derive from frictions in observing or acting on productivity in the labor marketplace (Acemoglu and Pischke 1998), where professions may provide a function of certifying membership and organizing labor. Professional membership may sometimes be exclusive, as in the case of physicians, but it may nonetheless be a highly imperfect proxy for productivity. When organizations cannot observe or act on individual productivity, they may need to rely on second-best policies that adhere to professional lines to improve efficiency. However, if individual productivity could be better ascertained and utilized, there may be scope for even larger efficiency gains for workplaces moving beyond the instrument of the professional class.

References

- Abadie, Alberto.** 2003. “Semiparametric Instrumental Variable Estimation of Treatment Response Models.” *Journal of Econometrics* 113 (2): 231–263.
- Abaluck, Jason, Leila Agha, Chris Kabrhel, Ali Raja, and Arjun Venkatesh.** 2016. “The Determinants of Productivity in Medical Testing: Intensity and Allocation of Care.” *American Economic Review* 106 (12): 3730–3764.
- Abbott, Andrew.** 2014. *The System of Professions: An Essay on the Division of Expert Labor*. Chicago: University of Chicago Press.
- Acemoglu, Daron, and David Autor.** 2011. “Skills, Tasks and Technologies: Implications for Employment and Earnings.” In *Handbook of Labor Economics*, edited by David Card and Orley Ashenfelter, vol. 4B, 1043–1171. Amsterdam: Elsevier.
- Acemoglu, Daron, and Jörn-Steffen Pischke.** 1998. “Why Do Firms Train? Theory and Evidence.” *Quarterly Journal of Economics* 113 (1): 79–119.
- Adhvaryu, Achyuta, Namrata Kala, and Anant Nyshadham.** 2022. “Management and Shocks to Worker Productivity.” *Journal of Political Economy* 130 (1): 1–47.
- Agency for Healthcare Research and Quality.** 2021. “Prevention Quality Indicators Technical Specifications Updates.” Accessed October 12, 2022. https://qualityindicators.ahrq.gov/archive/pqi_techspec/icd10_v2021.

- Alexander, Diane, and Molly Schnell.** 2019. “Just What the Nurse Practitioner Ordered: Independent Prescriptive Authority and Population Mental Health.” *Journal of Health Economics* 66:145–162.
- 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): 151–184.
- American Association of Nurse Practitioners.** 2020. “The Path to Becoming a Nurse Practitioner (NP).” Accessed February 23, 2022. <https://www.aanp.org/news-feed/explore-the-variety-of-career-paths-for-nurse-practitioners>.
- . 2021. “State Practice Environment.” Accessed February 23, 2022. <https://www.aanp.org/advocacy/state/state-practice-environment>.
- American College of Emergency Physicians.** 2016. “Emergency Department Crowding: High Impact Solutions.” Accessed February 20, 2022. https://www.acep.org/globalassets/sites/acep/media/crowding/empc_crowding-ip_092016.pdf.
- Anderson, D. Mark, Ryan Brown, Kerwin Kofi Charles, and Daniel I. Rees.** 2020. “Occupational Licensing and Maternal Health: Evidence from Early Midwifery Laws.” *Journal of Political Economy* 128 (11): 4337–4383.
- Andrews, Michelle, and McKenzie Beard.** 2024. “Like Doctors, More Nurse Practitioners Are Heading into Specialty Care.” *Washington Post*, June 17. Accessed June 17, 2024. <https://www.washingtonpost.com/politics/2024/06/17/like-doctors-more-nurse-practitioners-are-leaving-primary-care-behind/>.
- Ashley, Jo Ann.** 1976. *Hospitals, Paternalism, and the Role of the Nurse*. New York: Teachers College Press.
- Association of American Medical Colleges.** 2020. “The Road to Becoming a Doctor.” Accessed February 25, 2022. <https://www.aamc.org/system/files/2020-11/aamc-road-to-becoming-doctor-2020.pdf>.
- . 2021. *Physician Specialty Data Report*. Accessed April 25, 2024. <https://www.aamc.org/data-reports/workforce/data/number-people-active-physician-specialty-2021>.
- Beck, Andrew H.** 2004. “The Flexner Report and the Standardization of American Medical Education.” *JAMA* 291 (17): 2139–2140.

- Benkard, Lanier.** 2000. “Learning and Forgetting: The Dynamics of Aircraft Production.” *American Economic Review* 90 (4): 1034–1054.
- Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan.** 2004. “How Much Should We Trust Differences-in-Differences Estimates?” *Quarterly Journal of Economics* 119 (1): 249–275.
- Bhuller, Manudeep, Gordon B. Dahl, Katrine V. Løken, and Magne Mogstad.** 2020. “Incarceration, Recidivism, and Employment.” *Journal of Political Economy* 128 (4): 1269–1324.
- Brown, E. Richard.** 1979. *Rockefeller Medicine Men: Medicine and Capitalism in America*. Berkeley: University of California Press.
- Bureau of Labor Statistics.** 2021a. “Occupational Employment and Wages, May 2020 (Nurse Practitioners).” Accessed February 16, 2024. <https://www.bls.gov/oes/2020/may/oes291171.htm>.
- . 2021b. “Occupational Employment and Wages, May 2020 (Physicians, All Other; and Ophthalmologists, Except Pediatric).” Accessed February 16, 2024. <https://www.bls.gov/oes/2020/may/oes291228.htm>.
- Butler, Christopher C., David Gillespie, Patrick White, Janine Bates, Rachel Lowe, Emma Thomas-Jones, Mandy Wootton, Kerenza Hood, Rhiannon Phillips, Hasse Melbye, et al.** 2019. “C-Reactive Protein Testing to Guide Antibiotic Prescribing for COPD Exacerbations.” *New England Journal of Medicine* 381 (2): 111–120.
- Cairns, Christopher, and Kai Kang.** 2022. “National Hospital Ambulatory Medical Care Survey: 2019 Emergency Department Summary Tables.” Accessed February 3, 2023. <https://dx.doi.org/10.15620/cdc:115748>.
- California Health and Human Services Agency.** 2025. “Rates of Preventable Hospitalizations for Selected Medical Conditions by County (LGHC Indicator).” Accessed April 20, 2025. <https://data.chhs.ca.gov/dataset/rates-of-preventable-hospitalizations-for-selected-medical-conditions-by-county>.
- Cameron, A. Colin, and Douglas L. Miller.** 2015. “A Practitioner’s Guide to Cluster-Robust Inference.” *Journal of Human Resources* 50 (2): 317–372.

- Carrell, Scott, and James West.** 2010. “Does Professor Quality Matter? Evidence from Random Assignment of Students to Professors.” Publisher: The University of Chicago Press, *Journal of Political Economy* 118 (3): 409–432.
- Chan, David C.** 2016. “Teamwork and Moral Hazard: Evidence from the Emergency Department.” *Journal of Political Economy* 124 (3): 734–770.
- Chan, David C., Matthew Gentzkow, and Chuan Yu.** 2022. “Selection with Variation in Diagnostic Skill: Evidence from Radiologists.” *Quarterly Journal of Economics* 137 (2): 729–783.
- Chan, David C., and Jonathan Gruber.** 2020. “Provider Discretion and Variation in Resource Allocation: The Case of Triage Decisions.” *AEA Papers and Proceedings* 110:279–283.
- Chen, Yiqun.** 2021. “Team-Specific Human Capital and Team Performance: Evidence from Doctors.” *American Economic Review* 111 (12): 3923–3962.
- Chetty, Raj, John N. Friedman, and Jonah E. Rockoff.** 2014. “Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates.” *American Economic Review* 104 (9): 2593–2632.
- Currie, Janet, and W. Bentley MacLeod.** 2017. “Diagnosing Expertise: Human Capital, Decision Making, and Performance Among Physicians.” *Journal of Labor Economics* 35 (1): 1–43.
- . 2020. “Understanding Doctor Decision Making: The Case of Depression Treatment.” *Econometrica* 88 (3): 847–878.
- Currie, Janet, and Jonathan Zhang.** 2025. “Doing More with Less: Predicting Primary Care Provider Effectiveness.” *Review of Economics and Statistics* 107 (2): 289–305.
- Dahl, Gordon B., Andreas Ravndal Kostøl, and Magne Mogstad.** 2014. “Family Welfare Cultures.” *Quarterly Journal of Economics* 129 (4): 1711–1752.
- Dillender, Marcus, Anthony T Lo Sasso, Brian J Phelan, and Michael R Richards.** 2022. “Occupational Licensing and the Healthcare Labor Market.” NBER Working Paper 29665.
- Doyle, Joseph J., Steven M. Ewer, and Todd H. Wagner.** 2010. “Returns to Physician Human Capital: Evidence from Patients Randomized to Physician Teams.” *Journal of Health Economics* 29 (6): 866–882.

- Efron, Bradley.** 2016. “Empirical Bayes Deconvolution Estimates.” *Biometrika* 103 (1): 1–20.
- 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.
- Epstein, Andrew, and Sean Nicholson.** 2009. “The Formation and Evolution of Physician Treatment Styles: An Application to Cesarean Sections.” *Journal of Health Economics* 28 (6): 1126–1140.
- Fairman, Julie.** 2009. *Making Room in the Clinic: Nurse Practitioners and the Evolution of Modern Health Care*. New Brunswick: Rutgers University Press.
- Farronato, Chiara, Andrey Fradkin, Bradley Larsen, and Erik Brynjolfsson.** 2020. “Consumer Protection in an Online World: An Analysis of Occupational Licensing.” NBER Working Paper 26601.
- Fleming-Dutra, Katherine E., Adam L. Hersh, Daniel J. Shapiro, Monina Bartoces, Eva A. Enns, Thomas M. File, Jonathan A. Finkelstein, Jeffrey S. Gerber, David Y. Hyun, Jeffrey A. Linder, et al.** 2016. “Prevalence of Inappropriate Antibiotic Prescriptions Among US Ambulatory Care Visits, 2010-2011.” *JAMA* 315 (17): 1864–1873.
- Flexner, Abraham.** 1910. *Medical Education in the United States and Canada: A Report to the Carnegie Foundation for the Advancement of Teaching*. New York: Carnegie Foundation for the Advancement of Teaching.
- Freidson, Eliot.** 1974. *Professional Dominance: The Social Structure of Medical Care*. New Brunswick: Transaction Publishers.
- Gennaioli, Nicola, Rafael La Porta, Florencio Lopez-de-Silanes, and Andrei Shleifer.** 2013. “Human Capital and Regional Development.” *Quarterly Journal of Economics* 128 (1): 105–164.
- Gottlieb, Joshua D., Maria Polyakova, Kevin Rinz, Hugh Shiplett, and Victoria Udalova.** 2020. “Who Values Human Capitalists’ Human Capital? Healthcare Spending and Physician Earnings.” CES Working Paper CES-20-23.
- Gowrisankaran, Gautam, Keith Joiner, and Pierre Thomas Léger.** 2023. “Physician Practice Style and Healthcare Costs: Evidence from Emergency Departments.” *Management Science* 69 (6): 3202–3219.

- Hallett, Christine E.,** ed. 2016. *The History of Nursing. Major Themes in Health and Social Welfare*. New York: Routledge.
- Health Economics Resource Center.** 2021. “HERC Inpatient Average Cost Data.” Accessed February 23, 2022. <https://www.herc.research.va.gov/include/page.asp?id=inpatient>.
- Health Resources and Services Administration.** 2010. *The Registered Nurse Population: Findings from the 2008 National Sample Survey of Registered Nurses*. Accessed April 5, 2024. <https://data.hrsa.gov/DataDownload/NSSRN/GeneralPUF08/rnsurveyfinal.pdf>.
- Huang, David T., Donald M. Yealy, Michael R. Filbin, Aaron M. Brown, Chung-Chou H. Chang, Yohei Doi, Michael W. Donnino, Jonathan Fine, Michael J. Fine, Michelle A. Fischer,** et al. 2018. “Procalcitonin-Guided Use of Antibiotics for Lower Respiratory Tract Infection.” *New England Journal of Medicine* 379 (3): 236–249.
- Imbens, Guido W., and Donald B. Rubin.** 1997. “Estimating Outcome Distributions for Compliers in Instrumental Variables Models.” *Review of Economic Studies* 64 (4): 555–574.
- Institute of Medicine.** 2006. *Hospital-Based Emergency Care: At the Breaking Point*. Washington, DC: National Academies Press.
- Jackson, C Kirabo, Jonah E Rockoff, and Douglas O Staiger.** 2014. “Teacher Effects and Teacher-Related Policies.” *Annual Review of Economics* 6 (1): 801–825.
- Kleiner, Morris M., and Alan B. Krueger.** 2013. “Analyzing the Extent and Influence of Occupational Licensing on the Labor Market.” *Journal of Labor Economics* 31 (2): S173–S202.
- Kleiner, Morris M., Allison Marier, Kyoung Won Park, and Coady Wing.** 2016. “Relaxing Occupational Licensing Requirements: Analyzing Wages and Prices for a Medical Service.” *Journal of Law and Economics* 59 (2): 261–291.
- Kleinpell, Ruth, Michelle L Cook, and Diane L Padden.** 2018. “American Association of Nurse Practitioners National Nurse Practitioner Sample Survey: Update on Acute Care Nurse Practitioner Practice.” *Journal of the American Association of Nurse Practitioners* 30 (3): 140–149.
- Kline, Patrick M., Evan K. Rose, and Christopher R. Walters.** 2022. “Systemic Discrimination Among Large U.S. Employers.” *Quarterly Journal of Economics* 137 (4): 1963–2036.


- Larson, Magali S.** 1979. *The Rise of Professionalism: A Sociological Analysis*. Berkeley: University of California Press.
- Laurant, Miranda, David Reeves, Rosella Hermens, Jose Braspenning, Richard Grol, and Bonnie Sibbald.** 2005. "Substitution of Doctors by Nurses in Primary Care." *Cochrane Database of Systematic Reviews* 2.
- Lazuka, Volha.** 2018. "The Long-Term Health Benefits of Receiving Treatment from Qualified Midwives at Birth." *Journal of Development Economics* 133:415–433.
- Liu, Chuan-Fen, Paul L Hebert, Jamie H Douglas, Emily L Neely, Christine A Sulc, Ashok Reddy, Anne E Sales, and Edwin S Wong.** 2020. "Outcomes of Primary Care Delivery by Nurse Practitioners: Utilization, Cost, and Quality of Care." *Health Services Research* 55 (2): 178–189.
- Lutfiyya, May Nawal, Lisa Tomai, Bianca Frogner, Frank Cerra, Daniel Zismer, and Stephen Parente.** 2017. "Does Primary Care Diabetes Management Provided to Medicare Patients Differ between Primary Care Physicians and Nurse Practitioners?" *Journal of Advanced Nursing* 73 (1): 240–252.
- Markovits, Daniel.** 2020. *The Meritocracy Trap: How America's Foundational Myth Feeds Inequality, Dismantles the Middle Class, and Devours the Elite*. New York: Penguin.
- McMichael, Benjamin J., and Sara Markowitz.** 2022. "Toward a Uniform Classification of Nurse Practitioner Scope of Practice Laws." *Medical Care Research and Review*.
- Milbank Memorial Fund.** 2024. "No One Can See You Now Five Reasons Why Access to Primary Care Is Getting Worse (and What Needs to Change)." Accessed June 4, 2024. https://www.milbank.org/wp-content/uploads/2024/02/Milbank-Scorecard-2024-APPENDIX_v02.pdf.
- Minni, Virginia.** 2023. "Making the Invisible Hand Visible: Managers and the Allocation of Workers to Jobs." Accessed March 9, 2024. https://www.dropbox.com/s/713rr6e0qez7rkr/Minni_JMP.pdf?e=1&dl=0.
- Morgan, Perri A., David H. Abbott, Rebecca B. McNeil, and Deborah A. Fisher.** 2012. "Characteristics of Primary Care Office Visits to Nurse Practitioners, Physician Assistants and Physicians in United States Veterans Health Administration Facilities, 2005 to 2010: A Retrospective Cross-Sectional Analysis." *Human Resources for Health* 10 (1): 1–8.

- National Sample Survey of Registered Nurses.** 2022. *2022 National Sample Survey of Registered Nurses*. Accessed May 25, 2024. <https://data.hrsa.gov/topics/health-workforce/nursing-workforce-survey-data#RegisteredNurses>.
- Neuman, Mark D., Brian T. Bateman, and Hannah Wunsch.** 2019. “Inappropriate Opioid Prescription After Surgery.” *The Lancet* 393 (10180): 1547–1557.
- New Jersey Department of Health.** 2011. “Health Care Quality Assessment.” Accessed May 16, 2025. <https://www.nj.gov/health/healthcarequality/documents/qi/pqi2011.pdf>.
- Nyweide, David J, Denise L Anthony, Julie PW Bynum, Robert L Strawderman, William B Weeks, Lawrence P Casalino, and Elliott S Fisher.** 2013. “Continuity of Care and the Risk of Preventable Hospitalization in Older Adults.” *JAMA Internal Medicine* 173 (20): 1879–1885.
- Patel, Kant, and Mark E. Rushefsky.** 2004. *The Politics of Public Health in the United States*. Armonk: M.E. Sharpe.
- Perloff, Jennifer, Catherine M DesRoches, and Peter Buerhaus.** 2016. “Comparing the Cost of Care Provided to Medicare Beneficiaries Assigned to Primary Care Nurse Practitioners and Physicians.” *Health Services Research* 51 (4): 1407–1423.
- Perry, John J.** 2009. “The Rise and Impact of Nurse Practitioners and Physician Assistants on their Own and Cross-Occupation Incomes.” *Contemporary Economic Policy* 27 (4): 491–511.
- Rivkin, Steven G, Eric A Hanushek, and John F Kain.** 2005. “Teachers, Schools, and Academic Achievement.” *Econometrica* 73 (2): 417–458.
- Shapiro, Carl.** 1986. “Investment, Moral Hazard, and Occupational Licensing.” *Review of Economic Studies* 53 (5): 843–862.
- Silver, David.** 2021. “Haste or Waste? Peer Pressure and Productivity in the Emergency Department.” *Review of Economic Studies* 88 (3): 1385–1417.
- Smith, Matthew, Danny Yagan, Owen Zidar, and Eric Zwick.** 2019. “Capitalists in the Twenty-First Century.” *Quarterly Journal of Economics* 134 (4): 1675–1745.


- Stange, Kevin.** 2014. “How Does Provider Supply and Regulation Influence Health Care Markets? Evidence from Nurse Practitioners and Physician Assistants.” *Journal of Health Economics* 33:1–27.
- Starr, Paul.** 1982. *The Social Transformation of American Medicine: The Rise of a Sovereign Profession and the Making of a Vast Industry*. New York: Basic Books.
- Traczynski, Jeffrey, and Victoria Udalova.** 2018. “Nurse Practitioner Independence, Health Care Utilization, and Health Outcomes.” *Journal of Health Economics* 58:90–109.

Figure 1: First Stage

A. Number of Physicians




B. Treated by NP




Notes: This figure represents a graphical illustration of the first-stage estimation. Panel A shows a binned scatter plot of the number of physicians on duty versus the number of NPs on duty. Panel B shows a binned scatter plot of whether the case is treated by an NP versus the number of NPs on duty. To construct these binned scatter plots, we first residualize both the y-axis and x-axis variables with respect to the baseline control vector (i.e., ED-by-year, ED-by-month, ED-by-day-of-the-week, and ED-by-hour-of-the-day indicators) and then add means back for ease of interpretation. The coefficients report the estimated slope of the best-fit line between the y-axis and x-axis variables (conditional on the baseline control vector), with standard errors clustered by provider reported in parentheses.

Figure 2: Balance in Patient Characteristics




Notes: This figure shows balance in patient characteristics across the number of NPs on duty, conditional on the baseline controls (i.e., ED-by-year, ED-by-month, ED-by-day-of-the-week, and ED-by-hour-of-the-day indicators). To construct these binned scatter plots, we first residualize both the y -axis and x -axis variables with respect to the baseline controls and then add means back for ease of interpretation. The middle number on the y -axis of each panel reports the mean of the sample; the top and bottom numbers report the mean plus and minus a half standard deviation, respectively (except for Panel C which caps the top number at 100 since the mean plus a half standard deviation is beyond the maximum possible). The coefficients report the estimated slope of the best-fit line between the y -axis and x -axis variables (conditional on the baseline controls), with standard errors clustered by provider reported in parentheses. Each panel also reports p -values for the coefficient estimates. For readability of the coefficients, Panels B, C, and D scale up the dependent variable by 100. Predicted 30-day mortality is generated from a linear regression of actual 30-day mortality on patient characteristics $\mathbf{X}_i$ in Equations (1) and (2), including demographics, comorbidities, prior health care use, vital signs, and three-digit primary diagnosis indicators.

Figure 3: Reduced Form and Balance




Notes: This figure shows binned scatter plots of patients' actual and predicted outcomes on the y-axis versus the number of NPs on duty on the x-axis, controlling for the baseline control vector (i.e., ED-by-year, ED-by-month, ED-by-day-of-the-week, and ED-by-hour-of-the-day indicators). Panel A reports results for log length of stay; Panel B reports results for log cost of the ED visit; Panel C reports results for 30-day preventable hospitalization. The solid circles and lines represent patients' actual outcomes. The hollow circles and dashed lines represent patients' predicted outcomes generated based on patient characteristics X_i included in Equations (1) and (2), including patient demographics, comorbidities, prior health care use, vital signs, and three-digit primary diagnosis indicators. The reduced-form coefficients are estimated using Equation (2), with patients' actual outcomes as the dependent variable; the balance-test coefficients are estimated by regressing patients' predicted outcomes on the number of NPs on duty, conditional on the baseline control vector. Standard errors clustered by provider are reported in parentheses.

Figure 4: Stability of OLS and IV Estimates




Notes: This figure shows the robustness of our OLS and IV estimates to the inclusion of different sets of patient controls. We divide patient observable characteristics into eight subsets: (i) five-year age-bin indicators; (ii) marital status; (iii) gender; (iv) race indicators; (v) indicators for 31 Elixhauser comorbidities; (vi) vital signs; (vii) prior health care use; and (viii) indicators for three-digit patient primary diagnosis of the visit. We then run separate regressions that control for each of the $2^8 = 256$ different combinations of patient covariates for each outcome. Each n on the x -axis indicates the number of covariate subsets included. For each n , we plot the maximum, mean, and minimum of the estimated coefficients for the effect of NPs using all possible combinations with n (out of eight) subsets of patient covariates. The connected triangles and circles show the means of the estimated coefficients from OLS and IV regressions, respectively. The dashed lines connect the maximum and minimum of the estimated IV coefficients. The dotted lines connect the maximum and minimum of the estimated OLS coefficients. The coefficients at the bottom of each panel show the IV and OLS estimates with the full set of patient controls, with standard errors clustered by provider reported in parentheses. Panel A reports results for log length of stay. Panel B reports results for log cost of the ED visit. Panel C reports results for 30-day preventable hospitalization.

Figure 5: Heterogeneous Effects by Case Complexity and Severity




Notes: This figure shows heterogeneous effects of NPs by case complexity and severity. Panels A, C, and E divide cases into quartiles by their total number of Elixhauser comorbidities, with higher quartiles indicating more complex cases. Panels B, D, and F divide cases by whether the condition severity measured by the 30-day mortality of cases with the same three-digit primary diagnosis is equal to or above the 95th percentile of the sample. The solid and dashed lines show complier potential outcomes if they were treated by NPs and physicians, respectively. We estimate complier potential outcomes under NPs by the IV regression in Equation (4). We estimate complier potential outcomes under physicians by an IV regression similar to Equation (4) but with a dependent variable of $y_i \times (NP_i - 1)$, i.e., the interaction between patient outcome and the indicator for being treated by an NP minus one. Panels A-B, C-D, and E-F report results for log length of stay, log ED cost, and inpatient admission in the ED visit, respectively. The IV estimates of the effects of NPs for each patient subgroup are reported in Table 5.

Figure 6: Balance in Physician Characteristics




Notes: These panels are graphical representations of the balance-test regression at the ED-day level of physician average characteristics (weighted by the number of cases treated by each physician) on the number of NPs on duty, conditional on ED-by-year, ED-by-month, and ED-by-day-of-the-week indicators. Coefficients from the regressions are reported in each panel, along with standard errors (shown in parentheses) and *p*-values. To construct the binned scatter plots, we first residualize both the *y*-axis variable (average characteristics of physicians on duty) and the *x*-axis variable (the number of NPs on duty) with respect to ED-by-year, ED-by-month, and ED-by-day-of-the-week indicators and then add means back to aid in interpretation. The middle number on the *y*-axis of each panel reports the mean of the sample; the top and bottom numbers report the mean plus and minus a half standard deviation, respectively. The physician characteristics reported in Panels A-E are, respectively, age, gender, value-added, practice style in terms of patient log length of stay, and practice style in terms of patient log cost of the ED visit. Construction details of value-added and practice style are described in Appendix A.3.

Figure 7: Balance in NP Characteristics



Notes: These panels are graphical representations of the balance-test regression at the ED-day level of NP average characteristics (weighted by the number of cases treated by each NP) on the number of NPs on duty, conditional on ED-by-year, ED-by-month, and ED-by-day-of-the-week indicators. Coefficients from the regressions are reported in each panel, along with standard errors (shown in parentheses) and *p*-values. To construct the binned scatter plots, we first residualize both the *y*-axis variable (average characteristics of NPs on duty) and the *x*-axis variable (the number of NPs on duty) with respect to ED-by-year, ED-by-month, and ED-by-day-of-the-week indicators and then add means back to aid in interpretation. The middle number on the *y*-axis of each panel reports the mean of the sample; the top and bottom numbers report the mean plus and minus a half standard deviation, respectively. The NP characteristics reported in Panels A-E are, respectively, age, gender, value-added, practice style in terms of patient log length of stay, and practice style in terms of patient log cost of the ED visit. Construction details of value-added and practice style are described in Appendix A.3.

Figure 8: Distribution of Provider Productivity, Patient Assignments, and Wages



Notes: Panel A reports the deconvolved distributions of provider effects on log total spending per case, with the provider effects on log total spending per case converted to total annual spending on the x -axis. See Appendix A.5.3 for details of the estimation. The solid and dashed lines show the deconvolved distributions for NPs and physicians, respectively. Panel B plots histograms of the average predicted 30-day mortality of patients assigned to each provider. Predicted 30-day mortality is winsorized at the value of 5 percentage points. Panel C plots histograms of provider wages observed in the VHA data (inflation-adjusted to the year 2020). Wages are winsorized at the value of \$450,000. In Panels B and C, the white and gray bins show results for NPs and physicians, respectively.

Table 1: Characteristics of Baseline Sample

	All	Treated by NP	Treated by physician	<i>p</i> -value
Age	62.05 [15.80]	60.72 [15.87]	62.46 [15.75]	0.00
Married	0.424 [0.494]	0.424 [0.494]	0.424 [0.494]	0.80
Male	0.905 [0.293]	0.904 [0.295]	0.906 [0.292]	0.00
Black	0.270 [0.444]	0.271 [0.445]	0.270 [0.444]	0.12
White	0.708 [0.455]	0.705 [0.456]	0.709 [0.454]	0.00
Asian/Pacific Islander	0.021 [0.142]	0.021 [0.144]	0.020 [0.142]	0.04
Outpatient visits in prior year	6.242 [7.284]	5.658 [6.361]	6.423 [7.538]	0.00
Inpatient stays in prior year	0.612 [1.543]	0.431 [1.249]	0.668 [1.620]	0.00
Elixhauser comorbidity count	3.599 [3.018]	3.190 [2.772]	3.726 [3.079]	0.00
Length of stay (minutes)	162.09 [172.48]	119.53 [131.28]	175.29 [181.38]	0.00
ED cost (\$, inflation-adjusted to 2020)	939 [1,331]	813 [1,010]	978 [1,413]	0.00
Inpatient admission (%)	16.62 [37.23]	7.87 [26.92]	19.34 [39.50]	0.00
30-day preventable hospitalization (%)	1.23 [11.04]	0.75 [8.60]	1.39 [11.69]	0.00
30-day mortality (%)	1.25 [11.10]	0.63 [7.91]	1.44 [11.91]	0.00
Observations	1,118,836	264,789	854,047	

Notes: Column 1 shows average characteristics of all cases in our analysis sample. Columns 2 and 3 show average characteristics of cases treated by NPs and physicians in the sample, respectively. Standard deviations are reported in brackets; *p*-values of *t*-tests for the equivalence of means between cases treated by NPs and by physicians are shown in the last column.

Table 2: Effect of NPs on Resource Use and Patient Outcomes

	Dependent variable				
	Log length of stay (1)	Log ED cost (2)	Admission (3)	30-day mortality (4)	30-day preventable hospitaliza- tion (5)
Panel A: OLS					
NP assignment	-0.168 (0.034)	-0.101 (0.023)	-2.982 (0.460)	-0.211 (0.032)	-0.067 (0.030)
Panel B: Reduced form					
Number of NPs	0.020 (0.008)	0.013 (0.005)	0.019 (0.108)	-0.021 (0.021)	0.047 (0.021)
Panel C: IV					
NP assignment	0.110 (0.045)	0.070 (0.030)	0.103 (0.585)	-0.116 (0.115)	0.252 (0.120)
Full controls	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	4.608	6.483	16.625	1.247	1.234
S.D. dep. var.	1.161	0.878	37.230	11.099	11.041
Observations	1,110,798	1,108,961	1,118,836	1,118,836	1,118,836

Notes: This table shows estimates of the NP effect on patient log length of stay, log cost of the ED visit, hospital admission in the ED visit, 30-day mortality, and 30-day preventable hospitalization in Columns 1-5, respectively. Panel A reports the OLS estimates; Panel B reports the reduced-form estimates; Panel C reports the IV estimates. Sample sizes in Columns 1 and 2 are smaller than that reported in Column 1 of Table 1 due to missing outcomes for a small number of cases. The set of full controls includes ED-by-year, ED-by-month, ED-by-day-of-the-week, and ED-by-hour-of-the-day indicators and patient characteristics that include five-year age-bin indicators, marital status, gender, race indicators (white, Black, and Asian/Pacific Islanders, with other racial categories omitted as the reference group), indicators for 31 Elixhauser comorbidities, prior health care use (the number of outpatient visits and the number of inpatient stays in VHA facilities in the prior 365 days), vital signs, and indicators for three-digit ICD-10 code of patient primary diagnosis of the visit. Standard errors clustered by provider are shown in parentheses.

Table 3: Heterogeneous Effects by Provider Experience (Standardized)

	Dependent variable				
	Log length of stay (1)	Log ED cost (2)	Admission (3)	30-day mortality (4)	30-day preventable hospitaliza- tion (5)
Panel A: Provider specific experience					
NP assignment	0.101 (0.044)	0.072 (0.030)	0.092 (0.579)	-0.115 (0.116)	0.255 (0.121)
NP assignment $\times$ experience	-0.058 (0.025)	-0.042 (0.019)	-0.504 (0.331)	-0.001 (0.041)	0.016 (0.030)
Experience	-0.001 (0.006)	0.006 (0.010)	0.238 (0.314)	-0.001 (0.018)	-0.016 (0.014)
Panel B: Provider general experience					
NP assignment	0.086 (0.043)	0.062 (0.029)	0.089 (0.608)	-0.100 (0.116)	0.255 (0.121)
NP assignment $\times$ experience	-0.103 (0.056)	-0.035 (0.033)	0.340 (1.048)	0.088 (0.093)	0.043 (0.068)
Experience	-0.036 (0.015)	-0.013 (0.011)	-0.719 (0.221)	-0.012 (0.023)	-0.048 (0.025)
Full controls	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	4.608	6.483	16.625	1.247	1.234
S.D. dep. var.	1.161	0.878	37.230	11.099	11.041
Observations	1,110,798	1,108,961	1,118,836	1,118,836	1,118,836

Notes: Panel A shows heterogeneous effects of NPs by provider-specific experience in the case's condition, measured as the number of cases with the same three-digit primary diagnosis as the current case the provider has treated from the start of the study period to the day before the current case's visit. Panel B shows heterogeneous effects of NPs by provider general experience, measured as the number of cases (regardless of their diagnoses) the provider has treated from the start of the study period to the day before the current case's visit. For ease of interpretation, both specific and general measures of experience are standardized to have a mean of zero and a standard deviation of one for NPs and physicians separately. The outcome variables in Columns 1-5 are, respectively, log length of stay, log cost of the ED visit, hospital admission in the ED visit, 30-day mortality, and 30-day preventable hospitalization. All estimations include the full set of controls described in the notes to Table 2. Standard errors clustered by provider are shown in parentheses.

Table 4: Clinical Decision-Making

	Dependent variable				
	Consult (1)	CT (2)	X-ray (3)	Opioid (4)	Antibiotic (5)
NP assignment	0.026 (0.009)	0.012 (0.007)	0.020 (0.009)	-0.018 (0.006)	0.040 (0.022)
Full controls	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	0.226	0.145	0.291	0.088	0.639
S.D. dep. var.	0.418	0.352	0.454	0.283	0.480
Observations	1,118,836	1,118,836	1,118,836	1,118,836	123,395

Notes: This table shows IV estimates of the effect of NPs on various measures of clinical decision-making. The outcomes in Columns 1-5 are whether the patient receives during the ED visit formal consults, CT scans, X-rays, opioid prescriptions, and antibiotic prescriptions, respectively. Since antibiotics generally only apply to patients with infections, Column 5 restricts the sample to patients with respiratory or genitourinary system infections, two common types of infections. All estimations include the full set of controls described in the notes to Table 2. Standard errors clustered by provider are shown in parentheses.

Table 5: Heterogeneous Effects by Patient Characteristics

	Dependent variable				
	Log length of stay (1)	Log ED cost (2)	Admission (3)	30-day mortality (4)	30-day preventable hospitaliza- tion (5)
Panel A: Elixhauser comorbidity count					
1st quartile	0.042 (0.045)	0.028 (0.032)	-0.077 (0.636)	-0.147 (0.120)	0.555 (0.119)
2nd quartile	0.063 (0.044)	0.071 (0.030)	0.291 (0.642)	-0.041 (0.126)	0.438 (0.132)
3rd quartile	0.117 (0.048)	0.082 (0.031)	-0.245 (0.761)	-0.099 (0.178)	0.250 (0.186)
4th quartile	0.281 (0.066)	0.122 (0.041)	0.435 (1.476)	-0.203 (0.340)	-0.513 (0.347)
Panel B: Diagnosis predicted 30-day mortality					
< 95th percentile	0.080 (0.044)	0.064 (0.029)	-0.768 (0.573)	-0.077 (0.110)	0.361 (0.118)
≥ 95th percentile	0.988 (0.239)	0.247 (0.115)	26.140 (7.829)	-1.253 (2.127)	-2.989 (1.492)
Panel C: Diagnosis category					
Sepsis	1.480 (0.609)	0.095 (0.329)	44.880 (24.114)	24.533 (15.169)	11.117 (7.961)
Stroke	1.863 (0.677)	0.651 (0.311)	72.609 (31.758)	3.373 (6.038)	-0.062 (2.379)
AMI	0.806 (0.562)	1.780 (0.655)	123.007 (62.695)	-11.517 (9.684)	-3.219 (7.593)
Heart failure	1.125 (0.292)	0.088 (0.177)	20.263 (8.921)	1.262 (3.011)	-11.469 (5.265)
Other	0.097 (0.045)	0.067 (0.029)	-0.343 (0.578)	-0.147 (0.112)	0.332 (0.122)
Full controls	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	4.608	6.483	16.625	1.247	1.234
S.D. dep. var.	1.161	0.878	37.230	11.099	11.041
Observations	1,110,798	1,108,961	1,118,836	1,118,836	1,118,836

Notes: This table shows heterogeneous effects of NPs by patient characteristics. Panel A divides cases into quartiles by their total number of Elixhauser comorbidities, with higher quartiles indicating more complex cases. Panel B divides cases by whether the condition severity measured by the 30-day mortality of cases with the same three-digit primary diagnosis is equal to or above the 95th percentile of the sample. Panel C divides cases by their condition. The outcomes in Columns 1-5 are, respectively, log length of stay, log cost of the ED visit, hospital admission in the ED visit, 30-day mortality, and 30-day preventable hospitalization. All estimations include the full set of controls described in the notes to Table 2. Standard errors clustered by provider are shown in parentheses. Appendix Table A.9 reports outcome means for each group of cases in this table.

Online Appendix

A.1 Diagnosis Coding: NPs versus Physicians

In this appendix, we examine whether NPs and physicians differ significantly in reporting three-digit ICD-10 diagnoses. All diagnoses in our data are coded in ICD-10 during our study period from January 2017 to January 2020. As OLS estimation is likely to be confounded by patient selection, we leverage IV regressions that instrument for whether a case is treated by an NP using the number of NPs on duty. Specifically, we first create indicators for each of the 836 different three-digit ICD-10 primary diagnosis categories in our data (including one for the missing category). Then, for each diagnosis indicator, we run a separate 2SLS regression as follows to estimate whether NPs and physicians are significantly different in reporting:

$$y_i = \delta \text{NP}_i + \mathbf{T}_i \eta + \varepsilon_i, \quad (\text{A.1})$$

$$\text{NP}_i = \lambda Z_i + \mathbf{T}_i \zeta + v_i, \quad (\text{A.2})$$

where, similar to Equations (1) and (2), NP_i indicates whether case i is treated by an NP and Z_i denotes the instrument (i.e., the number of NPs on duty between 8 a.m. and 6 p.m., our analysis time window, at the ED on the day case i visits). The vector $\mathbf{T}_i$ includes ED-by-year, ED-by-month, ED-by-day-of-the-week, and ED-by-hour-of-the-day indicators. The coefficient of interest is δ . As with the main specification, we cluster standard errors by provider.

Panel A of Appendix Figure A.14 plots the distribution of t -statistics for the estimated δ coefficients from the 836 separate regressions that use each three-digit diagnosis indicator as the outcome variable. The share of t -statistics indicating a p -value below or equal to 0.05 is only 0.07, close to the null hypothesis of no differential three-digit diagnosis coding between NPs and physicians (i.e., a share of 0.05 of t -statistics indicating a p -value less than or equal to 0.05). Both the Shapiro-Wilk normality test and the test for normality on the basis of skewness and kurtosis suggest that we cannot reject the null hypothesis that the t -statistics are normally distributed, at least at the 10 percent level. Panel B of Appendix Figure A.14 further plots t -statistics against the prevalence of the three-digit diagnosis among physicians, showing that NPs are not more likely to report diagnoses that are more (or less) common.¹

The pattern of similar three-digit diagnosis coding between NPs and physicians could arise for the relatively straightforward cases that are compliers. Additional consults and diagnostics (Section 4.3) may also help NPs reach the same three-digit diagnosis as physicians. It is also worth noting that VHA ED providers' reimbursements are independent of patient diagnoses, and NPs and physicians are unlikely to have differential financial incentives in diagnosis coding.

1. We measure the prevalence of the diagnosis as the share of cases with the diagnosis among cases treated by physicians on days without any NP, to restrict the potential influence of patient sorting between NPs and physicians.

A.2 Characterizing Compliers and Never-Takers

This appendix describes the estimation of characteristics of compliers and never-takers. Following the approach developed by Abadie (2003), we characterize compliers by δ estimated through the 2SLS model specified in Equations (A.1) and (A.2), replacing the outcome variable y_i with $x_i \times \text{NP}_i$, i.e., the interaction between each patient characteristic x_i and the indicator for being treated by an NP. Results are discussed in Section 4.4 and shown in Column 2 of Appendix Table A.11.

To estimate the characteristics of never-takers, we follow a method by Dahl, Kostøl, and Mogstad (2014). We first collapse the data to the ED-day level. We then residualize the share of cases treated by NPs by ED-by-year, ED-by-month, and ED-by-day-of-the-week indicators. We define never-takers as cases treated by physicians in ED-day cells with the residual share of cases treated by NPs at least as high as the 90th percentile of ED-day cells with at least one case treated by NPs. There are no always-takers in our setting since patients cannot be assigned to NPs on days when no NPs are on duty.

Column 4 of Appendix Table A.11 reports the average characteristics of never-takers. For each characteristic, we compute the mean of never-takers as well as the ratio between the mean and the overall sample mean. We estimate standard errors for the means by bootstrap, blocking observations by provider with 500 replications. In line with the notion that NPs treat healthier cases than physicians do, Appendix Table A.11 shows that never-takers are the riskiest, followed by the overall sample, and finally, compliers. For example, the total number of Elixhauser comorbidities among never-takers, the overall sample, and compliers is, respectively, 4.0, 3.6, and 3.3; the average predicted 30-day mortality among these three types of cases is, respectively, 1.7, 1.2, and 0.9 percentage points.

A.3 Provider Value-Added and Practice-Style Measures

This appendix describes our construction of measures of provider value-added and practice styles, used to examine the exclusion restriction in Section 4.5. We consider physician value-added as a measure of risk-adjusted mortality outcomes and form this measure using leave-out data. Specifically, for physician p on day d , we measure

$$A_{p,d} = \frac{\sum_{i \in \mathcal{I}_p} \mathbf{1}(d(i) \neq d, Z_i = 0) \tilde{Y}_i}{\sum_{i \in \mathcal{I}_p} \mathbf{1}(d(i) \neq d, Z_i = 0)}, \quad (\text{A.3})$$

where $\tilde{Y}_i$ is risk-adjusted 30-day mortality, or the difference between the patient's actual and predicted 30-day mortality. To address possible finite-sample bias, we exclude patient cases on day d .² We also leave out cases visiting on days with any NPs on duty, to mitigate the concern of patient sorting between NPs and physicians.

Still, since cases are not experimentally assigned among physicians, $A_{p,d}$ may reflect both a physician's effect on patient outcomes and systematic patient-physician sorting under imperfect risk adjustment. As one way to assess the degree of such potential biases, we investigate the robustness of physician value-added

2. Specifically, there may be ED-day level shocks that are correlated with both the number of NPs on duty and the set of patients treated by a specific physician; these shocks can be influential in estimations with a finite sample.

estimates to patients' predicted mortality constructed using different risk adjusters, analogous to the test of student sorting biases in the teacher value-added literature (e.g., Chetty, Friedman, and Rockoff 2014). If patient sorting is important, the estimated physician value-added will change meaningfully with the addition of risk adjusters. Otherwise, our estimates should remain stable. Appendix Figure A.15 shows that physician value-added estimates are stable regardless of patient risk adjusters. We compare physician value-added measures constructed using (i) the most parsimonious set of risk adjusters that includes only age-bin and three-digit primary diagnosis indicators, (ii) the less parsimonious set that adds non-age demographics (gender, race, and marital status), and (iii) the set that further adds dummies for 31 Elixhauser comorbidities, with the baseline physician value-added constructed using the full set of patient covariates (i.e., demographics, Elixhauser comorbidities, prior health care use, vital signs, and three-digit diagnosis indicators). The correlations between measures (i)–(iii) and the baseline measure are all above 0.99. Note that these risk adjusters are important predictors of patient 30-day mortality: They alone explain 7 percent of the variation in 30-day mortality, with an F -statistic of 88 for joint significance.

We consider physician practice styles as measures of physician-chosen inputs to care. Specifically, we define practice style measures by Equation (A.3), but instead set $\tilde{Y}_i$ as the difference between the patient's actual and predicted log length of stay or log cost of the ED visit. As with value-added, we show the robustness of practice style estimates to different patient risk adjusters in Appendix Figure A.15.

We construct similar measures of value-added and practice style for NPs. As with physicians, we show the robustness of these estimates to different patient risk adjusters. Appendix Figure A.15 shows that NP value-added and practice-style estimates are highly stable among those constructed using (i) the most parsimonious set of risk adjusters that includes only age-bin and three-digit primary diagnosis indicators, (ii) the less parsimonious set that adds non-age demographics (gender, race, and marital status), (iii) the set that additionally includes dummies for 31 Elixhauser comorbidities, and (iv) the full set that further adds detailed controls for vital signs and prior health care use.

A.4 Augmenting Provider Supply with NPs

In this appendix, we ask, holding fixed the number of cases arriving and the number of physicians on duty, how additional NPs may affect patient wait time and downstream outcomes. The empirical specification takes the following form:

$$y_i = \sum_{n=0}^N \delta_n \times \mathbf{1}(Z_i = n) + N_i^c \gamma_1 + N_i^p \gamma_2 + \mathbf{T}_i \eta + \mathbf{X}_i \beta + \varepsilon_i, \quad (\text{A.4})$$

where $\mathbf{1}(Z_i = n)$ is an indicator for $n \in \{0, \dots, 4\}$ NPs being on duty at the ED on the day case i visits,³ and N_i^c and N_i^p are, respectively, the number of cases arriving and the number of physicians on duty at the ED on the day case i visits. We apply Equation (A.4) to the following outcomes: (i) wait time (the time between arrival at the ED and assignment to a provider); (ii) length of stay (the time between assignment to

3. The maximum n is 6, but only a small share (0.8 percent) of ED-day cells has 5 or 6 NPs on duty (Appendix Figure A.1). We thus group $n = 5$ or 6 with $n = 4$.

a provider and discharge from the ED); and (iii) log total spending of the ED visit, which is the sum of the cost of ED care, hospital admission in the ED visit, and 30-day preventable hospitalization.

Appendix Figure A.16 presents the estimated trade-off based on Equation (A.4). Panel A shows the trade-off between wait time and length of stay. Each dot plots a pair of δ_n with the estimated δ_n for wait time on the y-axis and the estimated δ_n for length of stay on the x-axis. Panel B similarly plots the trade-off between wait time and total spending per case (we convert the estimated δ_n for log total spending to spending per case).

As implied in Figure A.16, decreasing wait time by 30 minutes per case by hiring additional NPs increases length of stay by 10 minutes per case (Panel A) and increases total medical spending by \$436 per case, not including the cost of additional NP wages (Panel B; inflation-adjusted to 2020 dollars).⁴ Including the wage costs of additional NPs brings this figure to about \$480 per case.⁵ That is, roughly nine-tenths of the additional spending to reduce wait time by hiring NPs comes from the lower productivity of NPs, while only one-tenth comes from additional NP wage costs.

However, the cost of reducing wait time by hiring additional NPs declines with case complexity. As implied in Appendix Figure A.17, among the lowest-complexity quartile cases, reducing wait time by 30 minutes per case by hiring additional NPs increases total medical spending by about \$280 per case (not including the wage cost of additional NPs). Though reducing wait time by 30 minutes increases length of stay by a slightly greater magnitude in the lowest-complexity quartile cases than in the average case (12 versus 10 minutes per case), the difference is not statistically significant.

A.5 Distribution of Provider Effects on Total Spending

In this appendix, we estimate the distribution of provider effects on log total spending associated with the ED visit. We start by identifying provider effects using a just-identified IV model. Next, we estimate the variance of provider effects, using a split-sample approach to account for the bias due to sampling error in the estimated provider effects. We then apply an empirical Bayes deconvolution method, adapted by Kline, Rose, and Walters (2022) from Efron (2016), to recover the underlying population distribution of provider effects. We present the receiver operating characteristic (ROC) curves of provider productivity, case assignments, and wage payments as characteristics that may distinguish NPs and physicians. Finally, we assess the extent to which provider productivity relates to cases assigned and wages paid across individual providers.

A.5.1 Estimating Provider Effects

We generate a measure of total spending associated with each ED visit, which is the sum of the three main components of costs examined in Section 4: ED costs, hospital admission, and 30-day preventable hospitalization (we multiply the latter two components by the average cost of a hospital stay, \$19,220).

4. Figure A.16 shows that a 33-minute decline in wait time increases length of stay by 11 minutes and total spending by \$472 per case (comparing the first and last dots in each panel). We rescale the trade-off to a 30-minute decline in wait time for easier interpretation.

5. For this wage cost estimation, we divide the yearly NP wage by the average number of cases per ED-year and then multiply this figure by the number of additional NPs required.

We then estimate provider effects on total spending associated with the ED visit. To mitigate the impact of extreme values, we take the log of the medical expenditures. To account for the possibility that the treating provider is endogenous, we instrument for indicators for treating providers with indicators for on-duty providers in the ED-day cell of the patient's visit. The empirical specification is a just-identified 2SLS model as follows:

$$y_i = \sum_j \theta_j \mathbf{1}_{\{j(i)=j\}} + \mathbf{T}_i \eta + \mathbf{X}_i \beta + \varepsilon_i, \quad (\text{A.5})$$

$$\mathbf{1}_{\{j(i)=j\}} = \sum_j \lambda_j \mathbf{1}_{\{j \in \mathcal{I}_i\}} + \mathbf{T}_i \zeta + \mathbf{X}_i \gamma + v_i, \quad (\text{A.6})$$

where $\mathbf{1}_{\{j(i)=j\}}$ is an indicator for whether case i is treated by provider j , and $\mathbf{1}_{\{j \in \mathcal{I}_i\}}$ is an indicator for whether provider j is on duty in the ED-day cell of case i 's visit with $\mathcal{I}_i$ representing the set of providers on duty in the ED-day cell. The coefficients of interest are θ_j , representing provider effects. Since θ_j is only identified relative to one another for providers within the same ED, we make the natural normalization that the case-weighted mean of θ_j is zero within each ED, using linear constraints in the 2SLS estimation to yield valid standard errors.

The F -statistics for the joint significance of on-duty provider indicators in the first-stage regressions, i.e., Equation (A.6), have shares of 0.99 and 0.68 above 10 and 100, respectively, suggesting that provider availability is strongly predictive of the treating provider.⁶ Appendix Figure A.18 shows that patient characteristics are well balanced across the average characteristics (age, gender, and practice style) of on-duty providers, conditional on the baseline controls $\mathbf{T}_i$.⁷ In addition, the F -statistics for the joint significance of on-duty provider indicators from regressions of patient predicted log length of stay and predicted log cost of the ED visit on on-duty provider indicators conditioning on $\mathbf{T}_i$ are 2.4 and 2.1, respectively. These are much smaller than the corresponding F -statistics using the actual log length of stay and log cost of the ED visit as the outcome—which are 10.3 and 9.9, respectively. These results support the assumption that the set of on-duty providers is conditionally independent of the set of patients arriving, thereby validating our instruments.

A.5.2 Estimating Variance of Provider Effects

We estimate the variance of provider effects, within each professional class, on log total spending associated with the ED visit. The estimated provider effects $\hat{\theta}_j$ from Appendix A.5.1 yield a case-weighted variance of 0.054 for NPs, and 0.064 for physicians (see Appendix Table A.22).⁸ However, these estimates are upward-biased, due to sampling error resulting from the fact that provider effects are estimated on a finite sample. To account for such biases, we leverage a split-sample approach similar to that employed in earlier

6. Since $\mathbf{1}_{\{j(i)=j\}}$ is always zero for patients outside of the ED in which a provider practices, we report F -statistics from the first-stage regression in Equation (A.6) using observations in each ED separately.

7. We compute case-weighted average characteristics of on-duty providers, with the index case left out. For the practice style examined in this balance test, to deal with the concern about patient sorting between NPs and physicians, we use provider effects on patient log length of stay and log cost of the visit estimated by the 2SLS model in Equations (A.5) and (A.6).

8. Since provider effects are normalized to have a mean of 0 in each ED, the variance is interpretable as the within-ED variance of provider effects.

studies (e.g., Silver 2021). Specifically, we randomly split each provider's patients within each day into two approximately equal-sized partitions. We then estimate the 2SLS model in Equations (A.5) and (A.6) using each partition separately, yielding two fixed effect estimates for each provider $\hat{\theta}_{j,a}$ and $\hat{\theta}_{j,b}$. Suppressing the j subscript for simplicity, we have

$$\hat{\theta}_q = \theta + e_q, q \in \{a, b\},$$

where q indicates partitions, and e_q is partition-specific sampling error, such that $\text{Cov}(\theta, e_b) = \text{Cov}(e_a, \theta) = 0$. The random split of patients within each provider-day makes plausible the assumption that e_a and e_b are uncorrelated, i.e., $\text{Cov}(e_a, e_b) = 0$. We can therefore compute the variance of provider effects as the covariance of $\hat{\theta}_a$ and $\hat{\theta}_b$:

$$\begin{aligned} \text{Cov}(\hat{\theta}_a, \hat{\theta}_b) &= \text{Cov}(\theta + e_a, \theta + e_b) \\ &= \text{Cov}(\theta, \theta) + \text{Cov}(\theta, e_b) + \text{Cov}(e_a, \theta) + \text{Cov}(e_a, e_b) \\ &= \text{Var}(\theta). \end{aligned}$$

We perform this calculation separately for NPs and physicians.

Appendix Table A.22 reports the case-weighted variance of provider effects from the split-sample approach. The variance for physicians is estimated at 0.045, about 70 percent of the calculated variance without accounting for the bias due to sampling error. The variance from the split-sample approach suggests that, on average, a one-standard-deviation costlier physician raises medical spending by 21 percent per case. For NPs, the split-sample variance estimate is 0.048, suggesting that a one-standard-deviation costlier NP raises spending by 22 percent per case.

A.5.3 The Population Distribution of Provider Effects

We now estimate the distribution of provider effects by applying a non-parametric empirical Bayes deconvolution approach adapted by Kline, Rose, and Walters (2022) from Efron (2016). This approach extracts a flexible estimate of the distribution of population provider effects using provider effects $\hat{\theta}_j$ and their standard errors s_j estimated in Equations (A.5) and (A.6). Assuming provider z -scores $z_j = \hat{\theta}_j/s_j$ are distributed as

$$z_j|c_j \sim \mathcal{N}(c_j, 1), c_j \sim G_c,$$

where $c_j = \theta_j/s_j$ (i.e., the population analogue of z_j), the procedure first applies the Efron (2016) deconvolution procedure to yield a distribution of provider z -scores $\hat{G}_c$ with density function $\hat{g}_c(\cdot)$. The Efron (2016) procedure estimates $\hat{G}_c$ by maximum likelihood of parameters that represent coefficients on a set of splines, with a regularization parameter to tamp down excursions from a flat prior.

Next, assuming that s_j is independent of c_j , we can derive an estimate of the distribution of provider effects $\hat{F}$, with density function $\hat{f}(\cdot)$ for each value θ :

$$\hat{f}(\theta) = \frac{1}{J} \sum_{j=1}^J \frac{1}{s_j} \hat{g}_c(\theta/s_j). \quad (\text{A.7})$$

Following Kline, Rose, and Walters (2022), we assess the independence of z_j and s_j by reporting regressions of z_j on s_j . To account for possible correlated estimation errors in z_j and s_j , we also present split-sample versions of these regressions that randomly split cases for each provider into two approximately equal-sized partitions and regress z -scores from one partition on standard errors from the other partition. The results are reported in Appendix Table A.23, which show no significant relationship between z_j and s_j , suggesting that independence between z -scores and standard errors is plausible.

We apply the empirical Bayes deconvolution estimator to NPs and physicians separately.⁹ As in Kline, Rose, and Walters (2022), we calibrate the regularization parameter in the maximum likelihood estimation so that the variance of the deconvolved distribution of provider effects matches the corresponding split-sample variance estimates reported in Appendix Table A.22. We demean both the physician and NP distributions to have a mean of zero and then shift the distribution of NPs to the right by 0.067, where 0.067 is the 2SLS estimate of the NP effect on the log total spending associated with the ED visit obtained by Equations (1) and (2). Panel A of Figure 8 displays the deconvolved density of provider effects for NPs and physicians. For interpretation, we rescale provider effects on log total spending per case on the x -axis into annual spending, as follows: (i) estimating the empirical distribution of log total spending per case, where the log spending is residualized with respect to ED and provider fixed effects, with means added back; (ii) for each provider effect from the deconvolution, adding the provider effect to the spending distribution in (i); (iii) exponentiating the individual values of the distribution in (ii) and integrating over the distribution to calculate the mean total spending per case for each provider effect; and (iv) multiplying the mean in (iii) for each provider effect by the average number of cases a provider sees per year observed in the data.

Finally, using the deconvolved density of NP and physician effects on log total spending per ED case, we estimate the probability that a randomly drawn NP is costlier than a randomly drawn physician by

$$\Pr(\theta_j > \theta_{j'} | j \in \mathcal{J}_{NP}, j' \in \mathcal{J}_{MD}) = \int_0^1 \hat{F}_{MD}(\theta) d\hat{F}_{NP}(\theta), \quad (\text{A.8})$$

where $\hat{F}_{MD}(x)$ and $\hat{F}_{NP}(x)$ are the deconvolved cumulative distribution functions of physician effects and NP effects, respectively, and $\mathcal{J}_{MD}$ and $\mathcal{J}_{NP}$ are the sets of providers who are physicians and NPs, respectively. We find the probability that a randomly drawn NP is costlier than a randomly drawn physician, in terms of the total spending associated with the ED visit defined above, is 62 percent. Put differently, the probability that a randomly drawn NP is less costly than a randomly drawn physician is as high as 38 percent. This statistic remains large when we adjust the deconvolved productivity distributions to account for possible differences in treatment effects between the overall population and compliers: When we assume the average treatment effect is as large as twice the LATE estimate (higher than the LATE among the highest-complexity quartile patients), the probability that NPs are less costly remains large at 28 percent.

9. To restrict the inclusion of noisy provider effects, our deconvolution excludes providers with fewer than 150 cases. We set the support of provider effects to $[\theta^5 - SD, \theta^{95} + SD]$, where θ^5 , θ^{95} , and SD are, respectively, the 5th percentile, 95th percentile, and standard deviation of estimated NP and physician effects for the NP and physician deconvolution, respectively.

A.5.4 ROC Curve Representation

The probability in Equation (A.8) is equivalent to the area under the curve (AUC) of an ROC curve. The ROC curve displays the performance of an exercise classifying providers by a certain characteristic. In the case of provider effects on log total spending, the AUC of 0.62 indicates relatively poor performance in classifying providers as NPs versus physicians, depending on their (true) provider effects from their respective deconvolved distribution.

We construct ROC curves based on respective provider characteristics of productivity, the riskiness of assigned patients (measured by predicted 30-day mortality), and wages, where we consider physicians as the “positive” class and NPs as the “negative” class. For each characteristic of productivity, assigned patient predicted mortality, and wages, a provider with a higher value of the characteristic is more likely to be a physician (i.e., in the positive class). We define productivity as the additive inverse of the provider effect on log total spending per ED visit: $\mu_j = -\theta_j$. For a given characteristic x , we plot the ROC curve with $1 - \hat{F}_{MD}^x$ (i.e., the true positive rate) on the y-axis and $1 - \hat{F}_{NP}^x$ (i.e., the false positive rate) on the x-axis, where $\hat{F}_{MD}^x$ and $\hat{F}_{NP}^x$ are the empirical cumulative distribution functions of x among physicians and NPs, respectively. For productivity, we use the deconvolved distributions previously described in Appendix A.5.3, noting that $\hat{F}_{MD}^\mu = 1 - \hat{F}_{MD}^\theta$ and $\hat{F}_{NP}^\mu = 1 - \hat{F}_{NP}^\theta$. Since we observe actual patient assignments and provider wages, we do not estimate the risks of assigned patients or provider wage payments in a regression framework, nor do we apply deconvolution to obtain a population prior distribution. Rather, we apply the empirical cumulative distribution functions of the average predicted 30-day mortality for patients assigned to each provider and each provider’s annualized full-time equivalent (“yearly”) wage (inflation-adjusted to 2020 dollars).

We show the ROC curves in Appendix Figure A.12. As mentioned above, the AUC based on productivity is 0.62. The AUCs based on assigned patient risks and wage payments are 0.75 and 0.99, respectively. Taken together, a provider’s assigned patient risks are qualitatively more predictive of professional class than is the provider’s productivity; the provider’s yearly wages are extremely predictive of professional class.

A.5.5 Case Assignments and Wage Payments versus Productivity

Separately for NPs and physicians, we assess the extent to which case assignments and wages relate to individual provider productivity. As before, we measure productivity as provider effects on log total spending per ED visit, with higher provider effects indicating lower productivity. The empirical specification takes the following form:

$$y_j = \alpha \tilde{\theta}_j + \mathbf{L}_j \pi + \varepsilon_j, \quad (\text{A.9})$$

where y_j is the average risk (measured by predicted 30-day mortality) of cases assigned to provider j or yearly wage of provider j (inflation-adjusted to 2020 dollars). $\mathbf{L}_j$ is a vector of ED indicators since provider effects are only identified relative to one another within EDs. Since provider effects $\hat{\theta}_j$ are estimated with

noise, we calculate empirical Bayes posteriors of provider effects as in Kline, Rose, and Walters (2022):

$$\tilde{\theta}_j = s_j \times \frac{\int x \varphi(z_j - x) \hat{g}_c(x) dx}{\int \varphi(z_j - x) \hat{g}_c(x) dx}, \quad (\text{A.10})$$

where φ denotes the standard normal density.

A.6 ED-Specific NP Effects

In this appendix, we estimate heterogeneity in the ED-specific NP effect. In separate 2SLS regressions for each ED ℓ , we estimate the NP effect using only cases at that ED:

$$\begin{aligned} y_i &= \delta_\ell \text{NP}_i + \mathbf{t}_i \eta_\ell + \mathbf{X}_i \beta_\ell + \varepsilon_i, \\ \text{NP}_i &= \lambda_\ell Z_i + \mathbf{t}_i \zeta_\ell + \mathbf{X}_i \gamma_\ell + v_i, \end{aligned}$$

where $\mathbf{t}_i$ is a vector of indicators for patient arrival year, month, day of the week, and hour of the day.

In Appendix Figure A.11, we plot the distribution of $\hat{\delta}_\ell$ for all EDs in our sample. We also plot the empirical Bayes posterior mean $\tilde{\delta}_\ell$ for each ED, calculated as

$$\tilde{\delta}_\ell = w_\ell \hat{\delta}_\ell + (1 - w_\ell) \hat{\delta}. \quad (\text{A.11})$$

The shrinkage factor is given by $w_\ell = \frac{\hat{\pi}^2}{s_\ell^2 + \hat{\pi}^2}$, where $\hat{\pi}^2$ and s_ℓ^2 are, respectively, the variance of the prior distribution of δ_ℓ and the variance of the sampling error for each $\hat{\delta}_\ell$. We calculate s_ℓ^2 as the square of the standard error of $\hat{\delta}_\ell$. We calculate $\hat{\pi}^2$ as the difference between the case-weighted variance of $\hat{\delta}_\ell$ and the case-weighted mean of s_ℓ^2 . Finally, $\hat{\delta}$ is the case-weighted mean of $\hat{\delta}_\ell$.

The gray bins in Appendix Figure A.11 plot the empirical Bayes posterior $\tilde{\delta}_\ell$ for each ED in our sample.¹⁰ The distribution of posteriors is more compressed than that of the raw estimates of ED-specific effects, reflecting shrinkage due to sampling error in the raw estimates. The results show a fair amount of heterogeneity. Nonetheless, most EDs exhibit positive effects of NPs on raising patient length of stay, cost of the ED visit, and 30-day preventable hospitalization rate.

A.7 Measuring Number of On-Duty Physicians

In this appendix, we describe the measure of on-duty physicians. If patients arriving during the analysis time window are riskier, providers on duty during that period may be more occupied, raising the likelihood that patients arriving late in the window are assigned to physicians starting shifts after 6 p.m., thereby increasing the observed number of physicians treating patients arriving during the analysis time window.

10. The figure reports results for all EDs in our main analysis sample for log length of stay and log cost (in total 44 such EDs). For 30-day preventable hospitalization, since it is relatively uncommon (occurs in less than 2 percent of the sample), the estimates are relatively imprecise when using observations from a specific ED; we thus include only EDs with at least 25,000 cases in the analysis sample (in total 20 such EDs).

As shown in Columns 1-2 of Appendix Table A.24, we find a positive correlation between the health risks of patients arriving during the analysis time window and the observed number of physicians treating these patients. Note, however, that the measured number of on-duty NPs is unlikely to be subject to the same endogeneity concerns as the measured number of on-duty physicians, since NPs hardly work outside of our analysis time window (only one percent of ED-day cells in our analyzed EDs and months have NPs treating patients arriving outside of 8 a.m.–6 p.m.). Further, as shown in Figure 2 and Appendix Figure A.4, patient characteristics are remarkably balanced across the number of NPs on duty.

To deal with the endogeneity in the observed number of on-duty physicians, we consider a measure excluding physicians who treat patients arriving at or after 10 p.m., to drop physicians starting shifts after 6 p.m. (Results are similar when using alternative cutoffs of 9 or 11 p.m.) As shown in Columns 3-4 of Appendix Table A.24, although the relationship between patient health risks and the number of on-duty physicians based on this adjusted physician measure is weaker than that based on the raw measure in Columns 1-2, there remains a statistically significant relationship between patients' number of comorbidities and the measured number of physicians.


To obtain a more quasi-experimental measure of on-duty physicians, we count them over a 24-hour window. That is, we include all physicians who treat patients arriving at any time during the day, rather than restricting to those treating patients arriving between 8 a.m. and 6 p.m. This measure would still be correlated with the total physician staffing during our analysis time window, and its broader time window can mitigate concerns about the relationship between patient health risks and the identified number of on-duty physicians. As shown in Columns 5-6 of Appendix Table A.24, patient health risks are balanced across values of this 24-hour physician count. Comparing the three measures of on-duty physicians described above, Appendix Table A.25 shows that, when patient risks are more balanced against a measure of physician availability, IV estimates using the NP share of on-duty providers as the instrument become closer to our main IV estimates in Table 2. In our robustness checks in Appendix Figure A.10 and Panel D of Appendix Table A.18, we count on-duty physicians using the 24-hour physician count, the most balanced measure with respect to patient risks.

In Appendix Table A.26, we report results using the NP share of physician equivalents, assuming different substitution rates between NPs and physicians. Both the first-stage and the reduced-form estimates decline with the assumed NP-physician substitution rate, which is intuitive: If the assumed substitution rate exceeds the true rate, the NP assignment probability will rise less with a unit increase in the estimated share of physician equivalents who are NPs than it does with the actual share; conversely, if the assumed substitution rate is below the true rate, the NP assignment probability will rise more with a unit increase in the estimated share of physician equivalents who are NPs than it does with the actual share. The more the assumed substitution rate is above (below) the true rate, the smaller (larger) the first-stage estimates can be relative to the truth. The reduced-form estimates also decline with the assumed substitution rate, given that a lower share of patients is assigned to NPs with a unit increase in the estimated share of physician equivalents who are NPs. However, the resulting IV estimates are stable.

References


- Abadie, Alberto.** 2003. “Semiparametric Instrumental Variable Estimation of Treatment Response Models.” *Journal of Econometrics* 113 (2): 231–263.
- Chetty, Raj, John N. Friedman, and Jonah E. Rockoff.** 2014. “Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates.” *American Economic Review* 104 (9): 2593–2632.
- Dahl, Gordon B., Andreas Ravndal Kostøl, and Magne Mogstad.** 2014. “Family Welfare Cultures.” *Quarterly Journal of Economics* 129 (4): 1711–1752.
- Efron, Bradley.** 2016. “Empirical Bayes Deconvolution Estimates.” *Biometrika* 103 (1): 1–20.
- Kline, Patrick M., Evan K. Rose, and Christopher R. Walters.** 2022. “Systemic Discrimination Among Large U.S. Employers.” *Quarterly Journal of Economics* 137 (4): 1963–2036.
- Silver, David.** 2021. “Haste or Waste? Peer Pressure and Productivity in the Emergency Department.” *Review of Economic Studies* 88 (3): 1385–1417.

Figure A.1: Number of NPs on Duty




Notes: This figure shows the histogram of the number of NPs on duty in an ED-day cell. The unit of observation is at the ED-day level.

Figure A.2: Number of NPs on Duty by Month and Day of the Week




Notes: This figure plots the share of ED-day cells with different numbers of NPs on duty in each month (Panel A) and on each day of the week (Panel B). Each color of the stacked bars represents a different number of NPs on duty, with the lightest color at the bottom representing ED-day cells with the lowest number of NPs on duty (i.e., zero) and the darkest color at the top representing ED-day cells with the highest number of NPs on duty (i.e., more than three). Since only a small share of ED-day cells have more than three NPs on duty, this figure groups ED-day cells with more than three NPs into the group “>3”.

Figure A.3: Standard Deviation of Number of NPs on Duty and NP Assignment




Notes: This figure shows standard deviations of the number of NPs on duty (solid dots) and whether the patient is assigned to an NP (hollow dots). The first dot shows the raw standard deviations without residualizing the number of on-duty NPs or NP assignment by any factors. The second through the last dots show the standard deviations of the residualized number of on-duty NPs and the residualized NP assignment. The following factors are added to the residualization stepwise (cumulatively) from the second to the last dot: (i) ED-by-year indicators; (ii) ED-by-month indicators; (iii) ED-by-day-of-the-week indicators; (iv) ED-by-hour-of-the-day indicators; (v) patient covariates $\mathbf{X}_i$ in Equations (1) and (2); (vi) patient volume in the ED-day cell; (vii) average wait time of patients in the ED-day cell (leaving out the index patient); and (viii) average characteristics of on-duty physicians (age, gender, value-added, and practice style; see Appendix A.3 for the construction of value-added and practice style). Average characteristics of on-duty NPs are not included in the residualization because by definition these characteristics are missing for the 38 percent of ED-day cells without any NPs on duty.

Figure A.4: Balance in Patient Characteristics




Notes: This figure shows estimated coefficients and 95 percent confidence intervals from regressions of each patient characteristic listed on the y-axis on the number of NPs on duty, controlling for the baseline control vector (i.e., ED-by-year, ED-by-month, ED-by-day-of-the-week, and ED-by-hour-of-the-day indicators). For improved readability, a few coefficients (and their confidence intervals) are scaled up and down by, e.g., 10, as shown by “x 10” and “/ 10” on the y-axis, respectively. At the bottom of the figure, we report the F -statistic from the joint F -test for all patient characteristics in a reverse regression with the number of NPs on duty as the dependent variable, conditioning on the baseline control vector. Standard errors are clustered by provider.

Figure A.5: Predicting Log Length of Stay and Log ED Cost




Notes: This figure shows estimated coefficients and 95 percent confidence intervals from regressions of patient log length of stay and log cost of the ED visit on patient characteristics, controlling for the baseline control vector (i.e., ED-by-year, ED-by-month, ED-by-day-of-the-week, and ED-by-hour-of-the-day indicators). For improved readability, a few coefficients (and their confidence intervals) are scaled up by 10, as shown by “x 10” on the y-axis. The bottom of each panel reports the F -statistic from the joint F -test of all patient characteristics, conditioning on the baseline control vector. Standard errors are clustered by provider.

Figure A.6: Monotonicity Test




Notes: This figure shows the first-stage regressions for cases of different characteristics. Panels A-F split the sample by, respectively, age (above versus below the median of the sample), marital status, gender, race (Black versus non-Black), total number of Elixhauser comorbidities (above versus below the median of the sample), and predicted 30-day mortality (above versus below the median of the sample). Predicted 30-day mortality is generated from a linear regression of actual 30-day mortality on patient characteristics $\mathbf{X}_i$ included in Equations (1) and (2), including patient demographics, comorbidities, prior health care use, vital signs, and three-digit primary diagnosis indicators. To construct these binned scatter plots, we residualize both the y-axis and x-axis variables with respect to the baseline control vector (i.e., ED-by-year, ED-by-month, ED-by-day-of-the-week, and ED-by-hour-of-the-day indicators) within each subsample and then add means back. The coefficients report the first-stage estimates for each subset of patients conditional on the baseline control vector, with standard errors clustered by provider reported in parentheses.

Figure A.7: Reduced Form and Balance




Notes: This figure shows binned scatter plots of patients' actual and predicted outcomes on the y-axis versus the number of NPs on duty on the x-axis, controlling for the baseline control vector (i.e., ED-by-year, ED-by-month, ED-by-day-of-the-week, and ED-by-hour-of-the-day indicators). Panels A and B report results for hospital admission in the ED visit and 30-day mortality, respectively. The solid circles and lines represent patients' actual outcomes. The hollow circles and dashed lines represent patients' predicted outcomes generated based on patient characteristics $\mathbf{X}_i$ included in Equations (1) and (2), including patient demographics, comorbidities, prior health care use, vital signs, and three-digit primary diagnosis indicators. The reduced-form coefficients are estimated using Equation (2), with patients' actual outcomes as the dependent variable; the balance-test coefficients are estimated by regressing patients' predicted outcomes on the number of NPs on duty, conditional on the baseline control vector. Standard errors clustered by provider are reported in parentheses.

Figure A.8: Stability of OLS and IV Estimates




Notes: This figure shows the robustness of our OLS and IV estimates to the inclusion of different sets of patient controls. We divide patient observable characteristics into eight subsets: (i) five-year age-bin indicators; (ii) marital status; (iii) gender; (iv) race indicators; (v) indicators for 31 Elixhauser comorbidities; (vi) vital signs; (vii) prior health care use; and (viii) indicators for three-digit patient primary diagnosis of the visit. We then run separate regressions that control for each of the $2^8 = 256$ different combinations of patient covariates for each outcome. Each n on the x -axis indicates the number of covariate subsets included. For each n , we plot the maximum, mean, and minimum of the estimated coefficients for the effect of NPs using all possible combinations with n (out of eight) subsets of patient covariates. The connected triangles and circles show the means of the estimated coefficients from OLS and IV regressions, respectively. The dashed lines connect the maximum and minimum of the estimated IV coefficients. The dotted lines connect the maximum and minimum of the estimated OLS coefficients. The coefficients at the bottom of each panel show the IV and OLS estimates with the full set of patient controls, with standard errors reported in parentheses. Panels A and B report results for hospital admission in the ED visit and 30-day mortality, respectively.

Figure A.9: Patient Assignment




Notes: Panels A-C report patient characteristics (age, total number of Elixhauser comorbidities, and predicted 30-day mortality, respectively) on the y-axis against the number of NPs on duty on the x-axis. The unit of observation is at the case level. Both the y-axis and x-axis variables are residualized with respect to the baseline control vector (ED-by-year, ED-by-month, ED-by-day-of-the-week, and ED-by-hour-of-the-day indicators), with sample means added back to aid in interpretation. The circles, triangles, and squares show binned scatter plots for all cases, cases treated by physicians, and cases treated by NPs, respectively. Panel D plots whether the case is treated by an NP against the number of cases arriving in the analysis time window (i.e., 8 a.m. to 6 p.m.) of the ED-day cell of the case's visit. Both the y-axis and x-axis variables are residualized with respect to the baseline control vector, with means added back.

Figure A.10: Controlling for Physician Equivalents




Notes: This figure reports IV estimates along with 90 percent confidence intervals based on Equations (1) and (2) with an additional control: physician equivalents. We calculate physician equivalents as the sum of the number of on-duty physicians and the number of on-duty NPs multiplied by the assumed NP-physician substitution rate. See Section 4.5 for details. Each panel reports IV estimates assuming that one NP substitutes for 0.25, 0.5, 0.75, and 1 physician, respectively, from left to right as noted on the x -axis. To address endogeneity in the observed number of on-duty physicians, we measure on-duty physicians as those treating patients arriving at any time in the ED-day cell of the case's visit, rather than only those treating patients arriving during the analysis time window (8 a.m.–6 p.m.), as discussed in Appendix A.7.

Figure A.11: ED-Specific Estimates of NP Effect




Notes: This figure reports the distribution of ED-specific IV estimates of the NP effect. Panels A, B, and C report results for the NP effect on log length of stay, log cost of the ED visit, and 30-day preventable hospitalization, respectively. The white bins show the histogram of ED-specific IV estimates without any adjustment to account for estimation noise. The gray bins show the histogram of ED-specific IV estimates with empirical Bayes adjustments (see details in Appendix A.6). The dashed lines show the standard normal density with a variance of the prior distribution of ED-specific IV estimates for each outcome.

Figure A.12: Receiver Operating Characteristic (ROC) Curves



Notes: This figure displays the ROC curves for provider productivity (in the solid line), predicted 30-day mortality of cases assigned to the provider (in the dot-dashed line), and provider wages (in the dashed line). The dotted line plots the 45-degree line. Physicians are defined as the “positive” class, and NPs are defined as the “negative” class. The ROC curves are computed based on the distributions displayed in Figure 8. See Appendix A.5.4 for details.


Figure A.13: Case Assignment and Wage Payment versus Productivity




Notes: Panel A shows binned scatter plots of average risks (measured by predicted 30-day mortality) of cases assigned to a provider on the y-axis versus posterior mean provider effects on log total cost per ED visit on the x-axis. Panel B shows binned scatter plots of provider yearly wage on the y-axis versus posterior mean provider effects on log total cost per ED visit on the x-axis. Both the y-axis and x-axis variables are residualized with respect to ED indicators, with means added back for ease of interpretation. Wages are inflation-adjusted to the year 2020. Coefficients from regressions of wages on posterior mean provider effects controlling for ED indicators are reported, with standard errors clustered by ED shown in parentheses. The hollow circles report results for NPs; the solid circles report results for physicians. See more details in Appendix A.5.5.

Figure A.14: Diagnosis Coding: NPs versus Physicians

A. Distribution of t -Statistics




B. t -Statistics versus Diagnosis Prevalence




Notes: Panel A plots the distribution of the t -statistics on whether NPs and physicians are significantly different in diagnosis coding from 836 separate regressions that use each three-digit diagnosis indicator as the outcome variable. The distribution is estimated using an Epanechnikov kernel with the optimal bandwidth and shown in the solid line. For comparison, the standard normal density is plotted in the dashed line. Panel B shows binned scatter plots of the t -statistics against the prevalence of the diagnosis (measured as the share of cases with the diagnosis among cases treated by physicians on days without any NPs, to restrict the influence of patient sorting between NPs and physicians). The coefficient from the regression of the t -statistics on prevalence is reported at the bottom of the panel, along with its standard error (shown in parentheses) and p -value.

Figure A.15: Stability of Provider Value-Added and Practice Style with Varying Patient Covariates




Notes: This figure shows the stability of provider value-added and practice style estimated using alternative patient covariates. Appendix A.3 provides details. The x-axis reports provider value-added/practice style constructed using the full set of patient covariates, including demographics (five-year age-bin indicators, marital status, gender, and race indicators), indicators for 31 Elixhauser comorbidities, prior health care use (the number of outpatient visits and the number of inpatient stays in VHA facilities in the prior 365 days), vital signs, and indicators for the three-digit primary diagnosis of the visit. The y-axis reports provider value-added/practice style constructed using alternative sets of patient covariates: parsimonious set 1 that includes demographics, three-digit primary diagnosis indicators, and 31 Elixhauser comorbidities; parsimonious set 2 that includes demographics and three-digit primary diagnosis indicators; and parsimonious set 3 that includes five-year age-bin and three-digit primary diagnosis indicators. Value-added and practice style constructed using parsimonious sets 1-3 are shown in squares, circles, and “+”, respectively, which largely overlap. Correlations 1-3 report correlations of value-added/practice style estimated using the full set of patient covariates with those using parsimonious sets 1-3, respectively. The solid lines show the 45-degree line. Panels A, C, and E report results for physicians. Panels B, D, and F report results for NPs.

Figure A.16: Wait Time versus Length of Stay and Total Spending




Notes: This figure shows changes in patient wait time, length of stay, and total spending with incremental numbers of NPs on duty, conditional on the number of physicians on duty, the number of cases arriving, and the controls included in the main specification (see details in Appendix A.4). Panel A presents the trade-off between wait time and length of stay. Panel B presents the trade-off between wait time and total spending of the ED visit, measured as the sum of the cost of ED care, hospital admission in the ED visit, and 30-day preventable hospitalization. The solid lines show the quadratic fit estimated based on the plotted points.

Figure A.17: Wait Time versus Length of Stay and Total Spending: Lowest-Complexity Quartile Cases



Notes: This figure shows, for cases in the lowest-complexity quartile, changes in patient wait time, length of stay, and total spending with incremental numbers of NPs on duty, conditional on the number of physicians on duty, the number of cases arriving, and the controls included in the main specification (see details in Appendix A.4). Panel A presents the trade-off between wait time and length of stay. Panel B presents the trade-off between wait time and total spending of the ED visit, measured as the sum of the cost of ED care, hospital admission in the ED visit, and 30-day preventable hospitalization. The solid lines show the quadratic fit estimated based on the plotted points.

Figure A.18: Balance of Patient Characteristics across On-Duty Provider Characteristics



Notes: This figure shows estimated coefficients and 95 percent confidence intervals from regressions of each patient characteristic listed on the y-axis on average characteristics of providers on duty in the ED-day cell of the patient's visit, controlling for baseline controls (ED-by-year, ED-by-month, ED-by-day-of-the-week, and ED-by-hour-of-the-day indicators). All average on-duty provider characteristics are case-weighted, with the index case left out. The average on-duty provider characteristics in each panel are, respectively, age, female share, practice style in terms of patient log length of stay, and practice style in terms of patient log cost of care at the ED. For readability, a few coefficients (and their confidence intervals) are scaled down by 10 and 20, as shown by "10" and "20" on the y-axis, respectively. The bottom of each panel reports the F -statistic from the joint F -test for all patient characteristics in a reverse regression with the average on-duty provider characteristic as the dependent variable, conditioning on the baseline controls. Standard errors are clustered by provider.

Table A.1: Characteristics of NPs and Physicians at VHA and Non-VHA

	VHA (ED)	Non-VHA (ED)	Non-VHA (all)
Panel A: NPs			
Female (%)	81.4	79.1	90.0
Age	51.3	42.9	44.8
Panel B: Physicians			
Female (%)	34.0	27.3	31.0
Age	48.1	45.8	50.4

Notes: Panel A reports summary statistics for NPs; Panel B reports summary statistics for physicians. Column 1 reports summary statistics for NPs/physicians working at the ED in our analysis sample. Column 2 reports summary statistics for NPs/physicians working at the ED observed in the 20 percent Medicare data (with age and gender information obtained from the Medicare Data on Provider Practice and Specialty (MD-PPAS)). To provide a description of providers outside of the ED, Column 3 reports characteristics of all NPs/physicians (regardless of whether they work at the ED or not) observed in the Medicare data. VHA ED NPs' and physicians' mean age and female share reported in Column 1 are slightly different from those reported in Figures 6 and 7 because Figures 6 and 7 weight means by the number of ED-day cells a provider works and patient volume.

Table A.2: Selection of Baseline Sample

Sample step	Description	Cases	Providers		
			NPs	Physicians	EDs
1. Build sample of ED cases from January 1, 2017, to January 31, 2020.	We restrict the sample to cases after the VHA directive granting NPs full practice authority in December 2016 and before COVID pandemic in the US.	7,886,164	547	5,749	146
2. Include only cases visiting during daytime.	Empirically NPs hardly work outside of the hours of 8 a.m. to 6 p.m. We drop cases visiting outside of the hours of 8 a.m. to 6 p.m. to focus on cases that could be assigned to an NP.	5,766,296	539	5,665	145
3. Include only cases in ED-month cells with NPs.	We restrict the sample to cases in ED-month cells that use NPs to treat patients.	3,597,347	521	3,781	111
4. Restrict EDs to those in which NPs and physicians are the only providers.	To focus attention on the margin between NPs and physicians and hold the population of cases seen by an NP or physician fixed, we drop EDs that use other provider types, mainly physician assistants.	1,119,396	156	1,348	44
5. Drop cases with missing demographics or extreme ages.	We drop cases with missing age or gender, or age above 99 or below 20.	1,118,836	156	1,348	44

Notes: This table reports changes in the sample size when applying each of the listed sample restrictions. Columns 3-6 report, respectively, the number of cases, NPs, physicians, and EDs remaining at each step.

Table A.3: EDs Included in Main Analysis versus All VHA EDs

	Included	All
Panel A: ED characteristics		
Daily census	30.69 [17.58]	36.06 [21.34]
Urban location	0.86 [0.35]	0.87 [0.34]
Panel B: Provider Characteristics		
Physician mean age	52.218 [11.288]	50.962 [10.979]
Share physicians female	0.286 [0.452]	0.326 [0.469]
NP mean age	51.245 [8.940]	50.915 [9.352]
Share NPs female	0.803 [0.398]	0.752 [0.432]
Panel C: Patient characteristics		
Age	61.987 [15.857]	61.770 [15.828]
Married	0.416 [0.493]	0.432 [0.495]
Male	0.905 [0.294]	0.901 [0.299]
Black	0.275 [0.447]	0.272 [0.445]
White	0.704 [0.456]	0.709 [0.454]
Asian/Pacific Islander	0.020 [0.140]	0.018 [0.133]
Outpatient visits in prior year	6.409 [7.697]	6.552 [7.717]
Inpatient stays in prior year	0.643 [1.612]	0.689 [1.684]
Elixhauser comorbidity count	3.626 [3.035]	3.705 [3.100]
Predicted 30-day mortality (%)	1.200 [2.753]	1.229 [2.782]

Notes: Columns 1 and 2 report the characteristics of EDs included in our main analysis sample and the characteristics of all VHA EDs, respectively. Panel A reports EDs' daily census in the analysis time window (i.e., 8 a.m. to 6 p.m.) and urban location. Panel B reports characteristics of the physicians and NPs in the EDs. Panel C reports characteristics of patients seen in the EDs. The standard deviation of each characteristic is reported in brackets. Patient and provider characteristics in the included EDs, i.e., Column 1, are slightly different from those reported in Table 1 and Figures 6 and 7 because Column 1 includes observations in the ED for all months between January 2017 and January 2020, regardless of whether the ED used NPs in the month.

Table A.4: Heterogeneous Effects by Provider Experience (Standardized, Cases in 2018-)

	Dependent variable				
	Log length of stay (1)	Log ED cost (2)	Admission (3)	30-day mortality (4)	30-day prevent. hosp. (5)
Panel A: Provider specific experience					
NP assignment	0.081 (0.046)	0.077 (0.032)	-0.268 (0.646)	-0.222 (0.142)	0.335 (0.149)
NP assignment $\times$ experience	-0.060 (0.025)	-0.050 (0.021)	-0.677 (0.308)	0.011 (0.044)	-0.012 (0.033)
Experience	-0.006 (0.005)	0.004 (0.009)	0.204 (0.303)	-0.008 (0.017)	-0.018 (0.012)
Panel B: Provider general experience					
NP assignment	0.104 (0.048)	0.089 (0.034)	-0.356 (0.700)	-0.238 (0.146)	0.347 (0.152)
NP assignment $\times$ experience	-0.130 (0.067)	-0.069 (0.040)	0.249 (1.337)	0.108 (0.114)	-0.029 (0.077)
Experience	-0.034 (0.015)	-0.009 (0.011)	-0.721 (0.230)	-0.017 (0.024)	-0.039 (0.025)
Full controls	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	4.637	6.529	16.304	1.251	1.226
S.D. dep. var.	1.133	0.887	36.940	11.114	11.005
Observations	742,968	741,027	747,510	747,510	747,510

Notes: This table reports heterogeneous effects of NPs by provider experience using cases visiting in 2018 or after. Panel A shows heterogeneity by provider specific experience in the case's condition, measured as the number of cases with the same three-digit primary diagnosis as the current case the provider has treated since the start of the study period to the day before the current case's visit. Panel B shows heterogeneity by provider general experience, measured as the number of cases (regardless of conditions) the provider has treated since the start of the study period to the day before the current case's visit. For ease of interpretation, both specific and general experience are standardized to have a mean of zero and a standard deviation of one for NPs and physicians separately. The outcome variables in Columns 1-5 are, respectively, log length of stay, log cost of the ED visit, hospital admission in the ED visit, 30-day mortality, and 30-day preventable hospitalization. All estimations include the full set of controls described in the notes to Table 2. Standard errors clustered by provider are shown in parentheses.

Table A.5: Heterogeneous Effects by Provider Experience (Standardized, Prior-Year)

	Dependent variable				
	Log length of stay (1)	Log ED cost (2)	Admission (3)	30-day mortality (4)	30-day prevent. hosp. (5)
Panel A: Provider specific experience					
NP assignment	0.078 (0.046)	0.074 (0.031)	-0.295 (0.641)	-0.221 (0.141)	0.333 (0.148)
NP assignment $\times$ experience	-0.053 (0.027)	-0.055 (0.023)	-0.646 (0.307)	0.016 (0.049)	-0.017 (0.035)
Experience	-0.009 (0.006)	0.009 (0.011)	0.260 (0.327)	-0.012 (0.018)	-0.012 (0.014)
Panel B: Provider general experience					
NP assignment	0.089 (0.046)	0.081 (0.033)	-0.331 (0.662)	-0.222 (0.143)	0.350 (0.150)
NP assignment $\times$ experience	-0.120 (0.088)	-0.063 (0.044)	-0.142 (1.162)	0.040 (0.091)	-0.117 (0.085)
Experience	-0.040 (0.017)	-0.002 (0.012)	-0.732 (0.235)	-0.008 (0.026)	-0.008 (0.026)
Full controls	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	4.637	6.529	16.304	1.251	1.226
S.D. dep. var.	1.133	0.887	36.940	11.114	11.005
Observations	742,968	741,027	747,510	747,510	747,510

Notes: This table reports heterogeneous effects of NPs by provider experience in the prior year. Panel A shows heterogeneity by provider specific experience in the case's condition, measured as the number of cases with the same three-digit primary diagnosis as the current case the provider has treated in the 365 days prior to the day of the current case's visit. Panel B shows heterogeneity by provider general experience, measured as the number of cases (regardless of conditions) the provider has treated in the 365 days prior to the day of the current case's visit. The sample is restricted to cases visiting in 2018 or after, to allow for at least a one-year look-back window for measuring experience in the prior 365 days. For ease of interpretation, both specific and general experience are standardized to have a mean of zero and a standard deviation of one for NPs and physicians separately. The outcome variables in Columns 1-5 are, respectively, log length of stay, log cost of the ED visit, hospital admission in the ED visit, 30-day mortality, and 30-day preventable hospitalization. All estimations include the full set of controls described in the notes to Table 2. Standard errors clustered by provider are shown in parentheses.

Table A.6: Heterogeneous Effects by Provider Experience (Standardized, Days)

	Dependent variable				
	Log length of stay (1)	Log ED cost (2)	Admission (3)	30-day mortality (4)	30-day prevent. hosp. (5)
NP assignment	0.105 (0.045)	0.065 (0.029)	0.225 (0.623)	-0.100 (0.117)	0.263 (0.123)
NP assignment $\times$ experience	-0.031 (0.068)	-0.036 (0.037)	1.098 (1.157)	0.126 (0.111)	0.101 (0.080)
Experience	-0.015 (0.014)	-0.002 (0.009)	-0.403 (0.227)	-0.022 (0.024)	-0.053 (0.026)
Full controls	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	4.608	6.483	16.625	1.247	1.234
S.D. dep. var.	1.161	0.878	37.230	11.099	11.041
Observations	1,110,798	1,108,961	1,118,836	1,118,836	1,118,836

Notes: This table reports heterogeneous effects of NPs by provider general experience measured by the number of days the provider has worked since the start of the study period to the day before the current case's visit. For ease of interpretation, the experience measure is standardized to have a mean of zero and a standard deviation of one for NPs and physicians separately. The outcome variables in Columns 1-5 are, respectively, log length of stay, log cost of the ED visit, hospital admission in the ED visit, 30-day mortality, and 30-day preventable hospitalization. All estimations include the full set of controls described in the notes to Table 2. Standard errors clustered by provider are shown in parentheses.

Table A.7: Opioid Use Disorder and Return Visit with Infection

	Dependent variable			
	Opioid use disorder		Infection return visit	
	30-day (1)	180-day (2)	30-day (3)	180-day (4)
NP assignment	-0.0005 (0.0015)	-0.0000 (0.0019)	0.0142 (0.0068)	0.0225 (0.0100)
Full controls	Yes	Yes	Yes	Yes
Mean dep. var.	0.022	0.037	0.068	0.187
S.D. dep. var.	0.146	0.190	0.251	0.390
Observations	1,118,836	1,118,836	123,395	123,395

Notes: Columns 1 and 2 report the IV estimates of the effect of NPs on the probability of opioid use disorder in 30 days and 180 days after the index ED visit, respectively. Columns 3 and 4 report the IV estimates of the effect of NPs on the probability of return visits with infection in 30 days and 180 days after the index ED visit, respectively. Columns 3 and 4 restrict the sample to patients with respiratory or genitourinary system infections, following the sample restriction for the antibiotic prescription analysis in Table 4. All estimations include the full set of controls described in the notes to Table 2. Standard errors clustered by provider are shown in parentheses.

Table A.8: Ten Most Common High-Mortality Diagnoses

ICD code	Description	30-day mortality (%)	Cases	Share
I50	Heart failure	5.56	12,637	0.221
N17	Acute kidney failure	6.47	4,278	0.075
R41	Other symptoms and signs involving cognitive functions and awareness	7.59	3,872	0.068
D64	Other anemias	5.01	3,634	0.064
I21	Acute myocardial infarction	7.50	3,162	0.055
A41	Other sepsis	11.51	2,754	0.048
J15	Bacterial pneumonia, not elsewhere classified	5.12	2,715	0.047
J96	Respiratory failure, not elsewhere classified	12.99	2,548	0.045
F03	Unspecified dementia	5.40	1,427	0.025
R62	Lack of expected normal physiological development in childhood and adults	16.55	1,033	0.018

Notes: This table summarizes the 10 most common three-digit diagnosis codes in the group of diagnoses with a 30-day mortality rate equal to or above the 95th percentile of the sample. The columns report, from the leftmost to the rightmost, the three-digit ICD-10 code, description of the code, 30-day mortality rate of cases with the diagnosis code, number of cases in the analysis sample with the diagnosis code, and share of cases with the diagnosis code among all cases with a three-digit diagnosis whose 30-day mortality is equal to or above the 95th percentile of the sample.

Table A.9: Outcomes of Patient Subgroups

	Log length of stay	Log ED cost	Admission	30-day mortality	30-day prevent. hosp.
Panel A: Elixhauser comorbidity count					
1st quartile	4.393 [1.145]	6.321 [0.871]	8.297 [27.583]	0.406 [6.359]	0.249 [4.981]
2nd quartile	4.517 [1.155]	6.424 [0.879]	12.526 [33.101]	0.705 [8.364]	0.538 [7.317]
3rd quartile	4.664 [1.158]	6.531 [0.875]	18.448 [38.787]	1.188 [10.833]	1.155 [10.683]
4th quartile	4.944 [1.110]	6.720 [0.832]	30.717 [46.132]	3.062 [17.227]	3.448 [18.245]
Panel B: Diagnosis predicted 30-day mortality					
< 95th percentile	4.563 [1.159]	6.452 [0.876]	14.157 [34.861]	0.853 [9.195]	1.038 [10.138]
≥ 95th percentile	5.452 [0.818]	7.059 [0.707]	62.423 [48.433]	8.572 [27.995]	4.869 [21.522]
Panel C: Diagnosis category					
Sepsis	5.724 [0.715]	7.260 [0.632]	89.587 [30.548]	11.526 [31.939]	2.442 [15.437]
Stroke	5.365 [1.027]	7.209 [0.707]	55.926 [49.658]	2.638 [16.031]	0.606 [7.760]
AMI	5.181 [0.968]	7.275 [0.929]	57.527 [49.438]	7.495 [26.336]	3.479 [18.327]
Heart failure	5.501 [0.708]	7.050 [0.566]	66.072 [47.348]	5.534 [22.866]	13.382 [34.047]
Other	4.592 [1.160]	6.470 [0.877]	15.659 [36.341]	1.150 [10.662]	1.084 [10.353]

Notes: This table shows the mean outcomes of each group of patients examined in Table 5. Standard deviations are reported in brackets. Panel A divides cases into quartiles by their total number of Elixhauser comorbidities, with higher quartiles indicating more complex cases. Panel B divides cases by whether the condition severity measured by the 30-day mortality of cases with the same three-digit ICD-10 primary diagnosis is equal to or above the 95th percentile of the sample. Panel C divides cases by their condition. The outcomes in Columns 1-5 are, respectively, log length of stay, log cost of the ED visit, hospital admission in the ED visit, 30-day mortality, and 30-day preventable hospitalization.

Table A.10: Robustness to Controlling for Average Risks of Patients Assigned to Physicians

	Dependent variable				
	Log length of stay (1)	Log ED cost (2)	Admission (3)	30-day mortality (4)	30-day prevent. hosp. (5)
Panel A: Baseline					
NP assignment	0.110 (0.045)	0.070 (0.030)	0.103 (0.585)	-0.116 (0.115)	0.252 (0.120)
Panel B: Control for average age					
NP assignment	0.107 (0.045)	0.070 (0.030)	0.101 (0.590)	-0.125 (0.116)	0.254 (0.122)
Panel C: Control for average Elixhauser comorbidity count					
NP assignment	0.106 (0.045)	0.069 (0.030)	0.117 (0.592)	-0.136 (0.116)	0.238 (0.122)
Panel D: Control for average predicted 30-day mortality					
NP assignment	0.100 (0.044)	0.068 (0.030)	0.191 (0.596)	-0.114 (0.117)	0.258 (0.124)
Full controls	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	4.608	6.483	16.625	1.247	1.234
S.D. dep. var.	1.161	0.878	37.230	11.099	11.041
Observations	1,110,798	1,108,961	1,118,836	1,118,836	1,118,836

Notes: This table shows the robustness of the estimates to controlling for the average health risks of cases (age, total number of Elixhauser comorbidities, and predicted 30-day mortality in Panels B, C, and D, respectively) assigned to physicians in the ED-day cell (leaving out the index case). For comparison, Panel A repeats our baseline IV estimates reported in Panel C of Table 2. The outcomes in Columns 1-5 are, respectively, log length of stay, log cost of the ED visit, hospital admission in the ED visit, 30-day mortality, and 30-day preventable hospitalization. All estimations include the full set of controls described in the notes to Table 2. Standard errors clustered by provider are reported in parentheses.

Table A.11: Complier and Never-Taker Characteristics

	All	Compliers		Never-takers	
	Mean	Mean	Ratio	Mean	Ratio
Age	62.05 (0.15)	61.11 (0.31)	0.98 [0.98 - 0.99]	63.69 (0.17)	1.03 [1.02 - 1.03]
Married	0.424 (0.004)	0.424 (0.008)	1.00 [0.97 - 1.04]	0.436 (0.007)	1.03 [1.00 - 1.06]
Male	0.905 (0.002)	0.905 (0.003)	1.00 [0.99 - 1.01]	0.917 (0.002)	1.01 [1.01 - 1.02]
Black	0.270 (0.011)	0.262 (0.019)	0.97 [0.83 - 1.11]	0.228 (0.015)	0.84 [0.73 - 0.95]
White	0.708 (0.011)	0.716 (0.019)	1.01 [0.96 - 1.06]	0.756 (0.015)	1.07 [1.03 - 1.11]
Asian/Pacific Islander	0.021 (0.001)	0.020 (0.002)	0.95 [0.74 - 1.15]	0.013 (0.001)	0.65 [0.55 - 0.76]
Outpatient visits in prior year	6.242 (0.080)	5.824 (0.129)	0.93 [0.89 - 0.97]	6.537 (0.110)	1.05 [1.01 - 1.08]
Inpatient stays in prior year	0.612 (0.014)	0.490 (0.029)	0.80 [0.71 - 0.89]	0.695 (0.026)	1.14 [1.05 - 1.22]
Elixhauser comorbidity count	3.599 (0.030)	3.324 (0.066)	0.92 [0.89 - 0.96]	3.965 (0.041)	1.10 [1.08 - 1.12]
Predicted 30-day mortality (%)	1.247 (0.032)	0.902 (0.067)	0.72 [0.62 - 0.83]	1.697 (0.049)	1.36 [1.28 - 1.44]

Notes: This table reports average characteristics for the overall sample, compliers, and never-takers. Complier characteristics are estimated by 2SLS regressions replacing the outcome variable y_i with $x_i \times \text{NP}_i$, i.e., the interaction between patient characteristic and the indicator for being treated by an NP, controlling for the baseline control vector (i.e., ED-by-year, ED-by-month, ED-by-day-of-the-week, and ED-by-hour-of-the-day indicators). Standard errors clustered by provider are reported in parentheses. Never-takers are defined as cases treated by physicians in ED-day cells with the residual share of cases treated by NPs at least as high as the 90th percentile of ED-day cells with at least one case treated by NPs. Residual shares are constructed by first collapsing the data to ED-day cells and then residualizing the share of cases treated by NPs by ED-by-year, ED-by-month, and ED-by-day-of-the-week indicators. Standard errors for the overall sample and never-takers are estimated by bootstrap, using 500 replications and blocking observations by provider. For each characteristic, the table reports the mean as well as the ratio between this mean and the overall sample mean. The 95 percent confidence interval for each ratio is shown in brackets. Predicted 30-day mortality is generated from a linear regression of actual 30-day mortality on patient characteristics $\mathbf{X}_i$ included in Equations (1) and (2), including patient demographics, comorbidities, prior health care use, vital signs, and three-digit primary diagnosis indicators.

Table A.12: Robustness to Controlling for On-duty Provider Practice Style

	Dependent variable				
	Log length of stay (1)	Log ED cost (2)	Admission (3)	30-day mortality (4)	30-day prevent. hosp. (5)
Panel A: Control for average practice styles of on-duty physicians					
NP assignment	0.091 (0.044)	0.057 (0.029)	0.013 (0.588)	-0.116 (0.116)	0.256 (0.122)
Panel B: Control for average practice styles of on-duty physicians and NPs					
NP assignment	0.091 (0.044)	0.059 (0.030)	0.021 (0.589)	-0.116 (0.116)	0.255 (0.122)
Full controls	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	4.608	6.483	16.625	1.247	1.234
S.D. dep. var.	1.161	0.878	37.230	11.099	11.041
Observations	1,110,798	1,108,961	1,118,836	1,118,836	1,118,836

Notes: Panel A reports the robustness of our IV estimates to controlling for on-duty physicians' average practice styles in terms of both log length of stay and log cost of the ED visit. Panel B reports the robustness of our IV estimates to controlling for on-duty physicians' and on-duty NPs' average practice styles in terms of both log length of stay and log cost of the ED visit. In Panel B, for ED-day cells without NPs, we replace the missing average NP practice styles with the mean practice styles of NPs in the analysis sample. The outcomes in Columns 1-5 are, respectively, log length of stay, log cost of the ED visit, hospital admission in the ED visit, 30-day mortality, and 30-day preventable hospitalization. All estimations include the full set of controls described in the notes to Table 2. Standard errors clustered by provider are reported in parentheses.

Table A.13: On-Duty Physician Value-Added and Outcomes of Patients Treated by NPs

	Dependent variable					
	Elixhauser comorbidity count (1)	Predicted 30-day mortality (2)	Log length of stay (3)	Log ED cost (4)	Admission (5)	30-day prevent. hosp. (7)
Physician value-added	-0.006 (0.013)	-0.021 (0.014)	-0.010 (0.008)	-0.005 (0.005)	-0.245 (0.199)	-0.023 (0.044)
Controls	Baseline	Baseline	Full	Full	Full	Full
Mean dep. var.	3.128	0.728	4.302	6.298	7.726	0.633
S.D. dep. var.	2.711	2.115	1.083	0.870	26.700	7.929
Observations	147,936	147,936	146,948	146,935	147,936	147,936

Notes: This table shows the balance in outcomes for cases treated by NPs across the average value-added of physicians on duty. See Appendix A.3 for construction details of physician value-added. The outcomes in Columns 1-7 are, respectively, the total number of Elixhauser comorbidities, predicted 30-day mortality, log length of stay, log cost of the ED visit, hospital admission in the ED visit, 30-day mortality, and 30-day preventable hospitalization. The sample is restricted to patients treated by NPs on days with one NP on duty and at least one physician on duty. Since Columns 1-2 examine the balance in patient characteristics, the set of controls includes only the baseline control vector (i.e., ED-by-year, ED-by-month, ED-by-day-of-the-week, and ED-by-hour-of-the-day indicators). The set of full controls in Columns 3-7 is detailed in the notes to Table 2. Standard errors clustered by provider are reported in parentheses.

Table A.14: NP Presence and Outcomes of Patients Treated by Physicians

	Dependent variable						
	Elixhauser comorbidity count (1)	Predicted 30-day mortality (2)	Log length of stay (3)	Log ED cost (4)	Admission (5)	30-day mortality (6)	30-day prevent. hosp. (7)
Panel A: Baseline results							
NPs on duty	0.012 (0.032)	-0.013 (0.029)	-0.002 (0.012)	0.001 (0.007)	0.358 (0.303)	-0.022 (0.098)	-0.116 (0.117)
Panel B: By tercile of case count in ED-day cell							
NPs on duty							
× Bottom two terciles	0.027 (0.043)	0.017 (0.038)	0.027 (0.016)	0.009 (0.010)	0.386 (0.402)	0.023 (0.142)	-0.084 (0.162)
× Top tercile	0.002 (0.039)	-0.035 (0.034)	-0.017 (0.014)	-0.004 (0.009)	0.340 (0.383)	-0.053 (0.114)	-0.134 (0.146)
Controls	Baseline	Baseline	Full	Full	Full	Full	Full
Mean dep. var.	3.535	1.070	4.545	6.486	15.386	1.051	1.324
S.D. dep. var.	2.996	2.758	1.276	0.836	36.081	10.200	11.432
Observations	68,863	68,863	68,214	68,208	68,863	68,863	68,863

Notes: This table shows the balance in outcomes of patients treated by physicians against the presence of NPs. The sample is restricted to patients arriving between 5 and 8 a.m. in ED-day cells with all patients arriving between 5 and 8 a.m. being assigned to physicians. Panel A shows baseline results. The empirical specification takes the form $y_i = \gamma \mathbf{1}(Z_i > 0) + \mathbf{T}_i \eta + \mathbf{X}_i \beta + \varepsilon_i$, where $\mathbf{1}(Z_i > 0)$ is an indicator for whether there are NPs on duty during 8 a.m.-12 p.m. of the ED-day cell of the patient's visit. Panel B shows heterogeneous effects by whether the total number of cases in the ED-day cell is in the top tercile of all ED-day cells. The empirical specification takes the form $y_i = \sum_{g=1}^G \mathbf{1}(\text{Group}_i = g) [\gamma_g \mathbf{1}(Z_i > 0) + \lambda_g] + \mathbf{T}_i \eta + \mathbf{X}_i \beta + \varepsilon_i$, where $\mathbf{1}(\text{Group}_i = g)$ is an indicator for whether the ED-day cell has a number of cases arriving between 5 and 8 a.m. in the top tercile or bottom two terciles of all ED-day cells. The outcomes in Columns 1-7 are, respectively, the total number of Elixhauser comorbidities, predicted 30-day mortality, log length of stay, log cost of the ED visit, hospital admission in the ED visit, 30-day mortality, and 30-day preventable hospitalization. Since Columns 1-2 examine balance in patient characteristics, the set of controls includes only the baseline control vector (i.e., ED-by-year, ED-by-month, ED-by-day-of-the-week, and ED-by-hour-of-the-day indicators). The set of full controls in Columns 3-7 is detailed in the notes to Table 2. Standard errors clustered by provider are reported in parentheses.

Table A.15: Robustness to Additional Controls

	Dependent variable				
	Log length of stay (1)	Log ED cost (2)	Admission (3)	30-day mortality (4)	30-day prevent. hosp. (5)
Panel A: Baseline					
NP assignment	0.110 (0.045)	0.070 (0.030)	0.103 (0.585)	-0.116 (0.115)	0.252 (0.120)
Panel B: Control for patient volume (linear controls)					
NP assignment	0.097 (0.045)	0.077 (0.030)	0.376 (0.595)	-0.105 (0.115)	0.244 (0.118)
Panel C: Control for patient volume (linear splines)					
NP assignment	0.095 (0.045)	0.078 (0.030)	0.381 (0.595)	-0.101 (0.115)	0.245 (0.118)
Panel D: Control for patient volume (restricted cubic splines)					
NP assignment	0.096 (0.045)	0.078 (0.030)	0.377 (0.595)	-0.101 (0.115)	0.245 (0.118)
Panel E: Control for patient volume (fixed effects)					
NP assignment	0.097 (0.045)	0.078 (0.030)	0.383 (0.595)	-0.100 (0.115)	0.241 (0.118)
Panel F: Control for wait time					
NP assignment	0.109 (0.045)	0.069 (0.030)	0.317 (0.594)	-0.074 (0.124)	0.250 (0.121)
Full controls	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	4.608	6.483	16.625	1.247	1.234
S.D. dep. var.	1.161	0.878	37.230	11.099	11.041
Observations	1,110,798	1,108,961	1,118,836	1,118,836	1,118,836

Notes: Panel A repeats the baseline results in Panel C of Table 2. Panels B, C, D, and E add controls for patient volume in the analysis time window (8 a.m. to 6 p.m.) of the ED-day cell of the patient's visit, with patient volume controlled linearly, as linear splines, as restricted cubic splines, and as fixed effects for five-case bins, respectively. Panel F adds a control for patient wait time. As wait time is potentially endogenous to patient health (healthier cases may be assigned a lower priority and hence wait longer), we add an instrument for wait time: the average wait time of other cases in the ED-day cell. The outcomes in Columns 1-5 are, respectively, log length of stay, log cost of the ED visit, hospital admission in the ED visit, 30-day mortality, and 30-day preventable hospitalization. All estimations include the full set of controls described under Table 2. Standard errors clustered by provider are reported in parentheses.

Table A.16: Patient-Provider Gender Match

	Dependent variable					
	Elixhauser comorbidity count (1)	Predicted 30-day mortality (2)	Log length of stay (3)	Log ED cost (4)	Admission (5)	30-day prevent. hosp. (7)
Female NP $\times$ male patient	0.007 (0.103)	-0.065 (0.100)	0.027 (0.017)	-0.012 (0.015)	0.355 (0.593)	-0.058 (0.093)
Female NP	0.029 (0.153)	0.062 (0.153)	0.006 (0.133)	0.191 (0.081)	1.327 (1.886)	-0.015 (0.091)
Male patient	0.743 (0.092)	0.745 (0.092)	-0.051 (0.014)	-0.016 (0.013)	0.262 (0.508)	0.103 (0.078)
Controls	Baseline	Baseline	Full	Full	Full	Full
Mean dep. var.	3.190	0.743	4.304	6.341	7.866	0.745
S.D. dep. var.	2.772	2.145	1.137	0.856	26.921	8.598
Observations	264,772	264,772	262,960	263,045	264,772	264,772

Notes: This table examines whether NPs treat patients of the opposite gender differently compared to the same gender. We restrict the sample to patients treated by NPs, and regress each outcome on the interaction between the indicator for female NPs and the indicator for male patients, the indicator for female NPs, and the indicator for male patients. Columns 1-2 examine the balance in patient characteristics and add controls for the baseline control vector (i.e., ED-by-year, ED-by-month, ED-by-day-of-the-week, and ED-by-hour-of-the-day indicators). Columns 3-7 add the full set of controls described in the notes to Table 2. The outcomes in Columns 1-7 are, respectively, the total number of Elixhauser comorbidities, predicted 30-day mortality, log length of stay, log cost of the ED visit, hospital admission in the ED visit, 30-day mortality, and 30-day preventable hospitalization. Standard errors clustered by provider are reported in parentheses.

Table A.17: Alternative Standard Error Clustering

	Dependent variable				
	Log length of stay (1)	Log ED cost (2)	Admission (3)	30-day mortality (4)	30-day prevent. hosp. (5)
Panel A: Baseline					
NP assignment	0.110 (0.045)	0.070 (0.030)	0.103 (0.585)	-0.116 (0.115)	0.252 (0.120)
Panel B: Clustering by ED-day cell					
NP assignment	0.110 (0.015)	0.070 (0.010)	0.103 (0.348)	-0.116 (0.113)	0.252 (0.112)
Panel C: Two-way clustering by ED-day cell and provider					
NP assignment	0.110 (0.045)	0.070 (0.030)	0.103 (0.581)	-0.116 (0.113)	0.252 (0.119)
Full controls	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	4.608	6.483	16.625	1.247	1.234
S.D. dep. var.	1.161	0.878	37.230	11.099	11.041
Observations	1,110,798	1,108,961	1,118,836	1,118,836	1,118,836

Notes: This table reports the robustness of our estimates to alternative standard error clustering approaches. Panel A repeats our baseline estimates that cluster standard errors by provider. Panel B clusters standard errors by ED-day cell. Panel C clusters standard errors using two-way clustering by ED-day cell and provider. The outcomes in Columns 1-5 are, respectively, log length of stay, log cost of the ED visit, hospital admission in the ED visit, 30-day mortality, and 30-day preventable hospitalization. All estimations include the full set of controls described in the notes to Table 2.

Table A.18: Alternative Instruments

	Dependent variable				
	Log length of stay (1)	Log ED cost (2)	Admission (3)	30-day mortality (4)	30-day prevent. hosp. (5)
Panel A: Baseline					
NP assignment	0.110 (0.045)	0.070 (0.030)	0.103 (0.585)	-0.116 (0.115)	0.252 (0.120)
Panel B: Include NPs with only one case					
NP assignment	0.102 (0.045)	0.076 (0.030)	-0.011 (0.594)	-0.127 (0.116)	0.285 (0.125)
Panel C: Leave out the index case					
NP assignment	0.123 (0.046)	0.074 (0.031)	0.198 (0.605)	-0.110 (0.121)	0.260 (0.126)
Panel D: Share of on-duty providers who are NPs					
NP assignment	0.084 (0.039)	0.052 (0.025)	-0.086 (0.459)	-0.094 (0.098)	0.254 (0.099)
Panel E: Indicator for any NPs on duty					
NP assignment	0.108 (0.049)	0.080 (0.030)	0.185 (0.643)	-0.048 (0.121)	0.211 (0.130)
Panel F: Leave-out share of cases assigned to NPs					
NP assignment	0.114 (0.050)	0.058 (0.031)	0.678 (0.591)	-0.043 (0.119)	0.234 (0.118)
Full controls	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	4.608	6.483	16.625	1.247	1.234
S.D. dep. var.	1.161	0.878	37.230	11.099	11.041
Observations	1,110,798	1,108,961	1,118,836	1,118,836	1,118,836

Notes: Panel A repeats the baseline results in Panel C of Table 2. Panel B reports the results using an alternative measure of the number of NPs on duty as the instrument, which includes NPs with only one case in the analysis time window of the ED-day cell. Panel C reports the results using the measure of on-duty NPs that leaves out the index case in counting the number of on-duty NPs as the instrument. Panel D uses the share of on-duty providers who are NPs in the ED-day cell (the number of on-duty NPs divided by the total number of on-duty providers in the ED-day cell) as the instrument. As discussed in Appendix A.7, we measure on-duty physicians as those treating patients arriving at any time in the ED-day cell of the case's visit. Panel E uses an indicator for any NPs on duty in the ED-day cell as the instrument. Panel F uses the share of cases assigned to NPs in the ED-day cell, leaving out the index case, as the instrument. As with the share of on-duty providers who are NPs, we estimate the leave-out share of cases assigned to NPs using cases arriving at any time in the ED-day cell. The outcomes in Columns 1-5 are, respectively, log length of stay, log cost of the ED visit, hospital admission in the ED visit, 30-day mortality, and 30-day preventable hospitalization. All estimations include the full set of controls described in the notes to Table 2. Standard errors clustered by provider are shown in parentheses.

Table A.19: Sample Restricted to ED-Day Cells with Zero or One NP

	Dependent variable				
	Log length of stay (1)	Log ED cost (2)	Admission (3)	30-day mortality (4)	30-day prevent. hosp. (5)
NP assignment	0.109 (0.052)	0.084 (0.032)	0.146 (0.686)	-0.042 (0.128)	0.219 (0.140)
Full controls	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	4.594	6.445	16.301	1.241	1.235
S.D. dep. var.	1.147	0.887	36.937	11.069	11.045
Observations	862,416	860,798	868,930	868,930	868,930

Notes: This table shows results using only patients in ED-day cells with zero or one NP on duty. The outcomes in Columns 1-5 are, respectively, log length of stay, log cost of the ED visit, hospital admission in the ED visit, 30-day mortality, and 30-day preventable hospitalization. All estimations include the full set of controls described in the notes to Table 2. Standard errors clustered by provider are shown in parentheses.

Table A.20: Including Hospital Admissions and Preventable Hospitalizations Outside VHA

	Dependent variable			
	Admission		30-day prevent. hosp.	
	VHA only (1)	VHA+Medicare (2)	VHA only (3)	VHA+Medicare (4)
NP assignment	0.806 (0.798)	0.806 (0.794)	0.485 (0.205)	0.371 (0.222)
Full controls	Yes	Yes	Yes	Yes
Mean dep. var.	20.054	20.267	1.704	2.166
S.D. dep. var.	40.040	40.199	12.941	14.556
Observations	545,791	545,791	543,253	543,253

Notes: This table shows the robustness of our results to including hospital admissions and 30-day preventable hospitalizations outside of the VHA by examining patients who enroll in both the VHA and traditional Medicare. The VHA provides linked Medicare claims for beneficiaries who are traditional Medicare enrollees. Columns 1 and 2 show the robustness of results for hospital admissions during the ED visit. Column 1 measures only hospital admissions in the VHA; Column 2 adds hospital admissions in the Medicare claims. To obtain full observation of hospital admissions in non-VHA hospitals, Columns 1 and 2 restrict the sample to patients who enroll in traditional Medicare in the month of the ED visit. Columns 3 and 4 show the robustness of results for 30-day preventable hospitalizations. Column 3 measures 30-day preventable hospitalizations in the VHA; Column 4 adds 30-day preventable hospitalizations in the Medicare claims. To obtain full observation of 30-day preventable hospitalizations in non-VHA hospitals, Columns 3 and 4 restrict the sample to patients who enroll in traditional Medicare in both the month of the ED visit and the month that follows. All estimations include the full set of controls described in the notes to Table 2. Standard errors clustered by provider are shown in parentheses.

Table A.21: Including All EDs Using NPs

	Dependent variable				
	Log length of stay (1)	Log ED cost (2)	Admission (3)	30-day mortality (4)	30-day prevent. hosp. (5)
NP assignment	0.107 (0.030)	0.047 (0.018)	-0.266 (0.393)	-0.087 (0.076)	0.141 (0.083)
Full controls	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	4.600	6.522	17.127	1.216	1.268
S.D. dep. var.	1.151	0.869	37.674	10.958	11.190
Observations	2,167,104	2,166,520	2,184,032	2,184,032	2,184,032

Notes: This table shows the results that include cases in all VHA EDs that use NPs regardless of whether the ED uses physician assistants (but exclude cases in ED-day cells with physician assistants to focus on the margin between NPs and physicians). This table expands the analysis sample from 1.1 million cases across 44 EDs in our main analysis to 2.2 million cases across 110 EDs. The outcomes in Columns 1-5 are, respectively, log length of stay, log cost of the ED visit, hospital admission in the ED visit, 30-day mortality, and 30-day preventable hospitalization. All estimations include the full set of controls described in the notes to Table 2. Standard errors clustered by provider are shown in parentheses.

Table A.22: Variance of Provider Effects on Medical Spending

	NPs	Physicians
Basic estimates	0.0537	0.0643
Split-sample estimates	0.0476	0.0445

Notes: This table reports the variance of provider effects on log total spending of the ED visit. Total spending of the ED visit is computed as the sum of the three main cost components: the cost of care at the ED, hospital admission, and 30-day preventable hospitalization (we multiply the latter two components by the average cost of a hospital stay, \$19,220). Row 1 reports variance of provider effects $\hat{\theta}_j$ estimated using Equations (A.5) and (A.6). To account for biases due to estimation noise in $\hat{\theta}_j$, Row 2 reports variance estimated using a split-sample approach (details in Appendix A.5.2). Column 1 reports the variance for NPs. Column 2 reports the variance for physicians.

Table A.23: Relationship between z -scores and Standard Errors

	Dependent variable: provider z -score			
	(1)	(2)	(3)	(4)
Provider std. error	0.166 (0.434)	0.203 (0.348)	-0.456 (0.385)	0.085 (0.471)
Estimation sample	Full	Full	Split	Split
Provider group	NPs	Physicians	NPs	Physicians
Mean dep. var.	0.018	-0.137	0.071	-0.056
S.D. dep. var.	1.332	1.681	1.227	1.537
Mean std. error	0.307	0.226	0.283	0.202
S.D. std. error	0.515	0.206	0.365	0.124
Providers	75	644	64	474
Observations	75	644	128	948

Notes: This table reports coefficients from regressions of provider-specific z -scores on associated standard errors. Columns 1 and 2 report results using z -scores and standard errors estimated based on the full sample. Columns 3 and 4 randomly split cases for each provider into two approximately equal-sized partitions and regress z -scores from one partition on standard errors from the other partition, stacking the two partitions in the regressions. Columns 1 and 3 report results for NPs. Columns 2 and 4 report results for physicians. The numbers of unique providers in Columns 1 and 2 are smaller than those reported in Section 2.3 because our deconvolution includes only providers with at least 150 cases (to restrict the inclusion of noisy provider effect estimates). Columns 3 and 4 have smaller numbers of providers than those in Columns 1 and 2, respectively, because Columns 3 and 4 further drop providers with fewer than 150 cases in each split sample.

Table A.24: Balance across Different Measures of Number of On-Duty Physicians

	Raw measure		Exclude cases since 10 p.m.		24-hour measure	
	Elixhauser comorbidity count (1)	Predicted 30-day mortality (2)	Elixhauser comorbidity count (3)	Predicted 30-day mortality (4)	Elixhauser comorbidity count (5)	Predicted 30-day mortality (6)
Number of physicians	0.030 (0.007)	0.025 (0.008)	0.020 (0.007)	0.010 (0.007)	0.006 (0.006)	0.007 (0.007)
Controls	Baseline	Baseline	Baseline	Baseline	Baseline	Baseline
Mean dep. var.	3.599	1.247	3.599	1.247	3.599	1.247
S.D. dep. var.	3.018	2.857	3.018	2.857	3.018	2.857
Observations	1,118,836	1,118,836	1,118,836	1,118,836	1,118,836	1,118,836

Notes: This table examines balance in patient health risks across different measures of the number of on-duty physicians. Columns 1-2 define the number of on-duty physicians as those treating patients arriving between 8 a.m. and 6 p.m. in the ED-day cell. Columns 3-4 define the number of on-duty physicians as those treating patients arriving between 8 a.m. and 6 p.m. but exclude physicians who also treated patients arriving at or after 10 p.m. in the ED-day cell. Columns 5-6 define the number of on-duty physicians as those treating patients arriving at any time in the ED-day cell. The odd-numbered columns examine balance in patients' number of Elixhauser comorbidities; the even-numbered columns examine balance in patients' predicted 30-day mortality. All estimations include the baseline control vector, i.e., ED-by-year, ED-by-month, ED-by-day-of-the-week, and ED-by-hour-of-the-day indicators. Standard errors clustered by provider are reported in parentheses.

Table A.25: Share of Providers Who are NPs as the Instrument: Different Measures of Number of On-Duty Physicians

	Dependent variable				
	Log length of stay (1)	Log ED cost (2)	Admission (3)	30-day mortality (4)	30-day prevent. hosp. (5)
Panel A: Number of physicians based on the raw measure					
NP assignment	0.041 (0.037)	0.047 (0.025)	-0.211 (0.450)	-0.041 (0.100)	0.218 (0.100)
Panel B: Number of physicians excluding physicians treating cases arriving ≥ 10 p.m.					
NP assignment	0.066 (0.038)	0.052 (0.025)	-0.159 (0.460)	-0.027 (0.103)	0.243 (0.102)
Panel C: Number of physicians based on the 24-hour measure					
NP assignment	0.084 (0.039)	0.052 (0.025)	-0.086 (0.459)	-0.094 (0.098)	0.254 (0.099)
Full controls	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	4.608	6.483	16.625	1.247	1.234
S.D. dep. var.	1.161	0.878	37.230	11.099	11.041
Observations	1,110,798	1,108,961	1,118,836	1,118,836	1,118,836

Notes: This table reports the IV estimates using the share of on-duty providers who are NPs, i.e., the number of on-duty NPs divided by the total number of on-duty NPs and physicians, as the instrument. Panel A measures the number of on-duty physicians as those treating patients arriving between 8 a.m. and 6 p.m. in the ED-day cell. Panel B measures the number of on-duty physicians as those treating patients arriving between 8 a.m. and 6 p.m. but excludes physicians who also treated patients arriving at or after 10 p.m. in the ED-day cell. Panel C measures the number of on-duty physicians as those treating patients arriving at any time in the ED-day cell. The outcomes in Columns 1-5 are, respectively, log length of stay, log cost of the ED visit, hospital admission in the ED visit, 30-day mortality, and 30-day preventable hospitalization. All estimations include the full set of controls described in the notes to Table 2. Standard errors clustered by provider are reported in parentheses.

Table A.26: Share of Physician Equivalents Who are NPs as the Instrument

	First stage	Reduced form				
	NP assignment (1)	Log length of stay (2)	Log ED cost (3)	Admission (4)	30-day mortality (5)	30-day prevent. hosp. (6)
Panel A: 1 NP = 0.25 physicians						
NP share	1.269 (0.106)	0.083 (0.048)	0.060 (0.031)	-0.477 (0.565)	-0.013 (0.146)	0.308 (0.138)
Panel B: 1 NP = 0.5 physicians						
NP share	1.064 (0.072)	0.067 (0.039)	0.051 (0.026)	-0.263 (0.477)	-0.024 (0.113)	0.263 (0.107)
Panel C: 1 NP = 0.75 physicians						
NP share	0.913 (0.064)	0.059 (0.033)	0.046 (0.022)	-0.175 (0.416)	-0.023 (0.095)	0.225 (0.090)
Panel D: 1 NP = 1 physician						
NP share	0.814 (0.059)	0.054 (0.030)	0.042 (0.020)	-0.129 (0.376)	-0.022 (0.084)	0.198 (0.080)
Full controls	Yes	Yes	Yes	Yes	Yes	Yes
Mean dep. var.	0.237	4.608	6.483	16.625	1.247	1.234
S.D. dep. var.	0.425	1.161	0.878	37.230	11.099	11.041
Observations	1,118,836	1,110,798	1,108,961	1,118,836	1,118,836	1,118,836

Notes: This table reports the first-stage and reduced-form estimates using the share of on-duty physician equivalents who are NPs, i.e., the on-duty physician equivalents who are NPs divided by the total physician equivalents on duty, as the instrument. Panels A-D report results assuming that an NP substitutes for 0.25, 0.5, 0.75, and 1 physician, respectively. The outcomes in Columns 1-6 are, respectively, whether the patient is assigned to an NP, log length of stay, log cost of the ED visit, hospital admission in the ED visit, 30-day mortality, and 30-day preventable hospitalization. All estimations include the full set of controls described in the notes to Table 2. Standard errors clustered by provider are reported in parentheses.