Journal of Economic Behavior and Organization 212 (2023) 689–707

Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

journal homepage: www.elsevier.com/locate/jebo

The effects of doctor strikes on patient outcomes: Evidence from the English NHS☆

George Stoye*, Max Warner

Institute for Fiscal Studies, 7 Ridgmount Street, London, WC1E 7AE, UK

A R T I C L E   I N F O

Article history:
Received 14 December 2022
Revised 31 May 2023
Accepted 11 June 2023

JEL classification:
I10
I14
J52

Keywords:
Doctor strikes
Medical strikes
Junior doctors
Healthcare inequalities

A B S T R A C T

Doctor strikes cause major disruption for hospitals and patients. Past attempts to estimate impacts on patients suffer from selection issues due to changing patient composition during strikes. We address these issues by exploiting differential hospital exposure to a 2016 'junior' doctor strike in England to estimate the impact of doctor strikes on patient outcomes. Using the pre-strike junior-senior doctor ratio to measure exposure, we show increased strike exposure led to larger reductions in elective volumes, but did not affect volumes, average mortality or readmission rates for emergency patients. However, greater exposure to the strike did lead to higher readmission rates for black emergency patients. This suggests that while hospitals managed to mitigate many of the negative effects of the strikes, disruptions from the strikes still had negative consequences for some minority groups.

1. Introduction

Doctor strikes are highly controversial events that affect many health systems around the world (Chima, 2020). When considering strike action and how to respond to it, unions, healthcare providers and governments must trade-off the short-term disruption to patients with potential longer-run benefits from changes in employment conditions. This potential harm to patients is often used to argue against strikes, and in many countries strikes by healthcare workers are either banned or highly regulated events (European Foundation for the Improvement of Living and Working Conditions, 2011).[1] However, there is currently little empirical evidence on the causal impact of such strikes on patient health.

☆ We thank Laura Schechter, two anonymous referees, Eric French, Carol Propper, Imran Rasul and Marcos Vera-Hernández for useful comments, as well as seminar participants at the IFS, the American European Health Economics Study Group 2022 and the European Association of Labour Economists Conference 2022. The authors thank NHS Digital for access to the Hospital Episode Statistics under data sharing agreement CON-205762-B8S7B, and the Department of Health and Social Care for access to the Electronic Staff Record (ESR). Figures published using the ESR may be different from the official workforce statistics published by NHS Digital. The views expressed are those of the authors and are not necessarily those of the Department of Health and Social Care. Funding: the authors gratefully acknowledge financial support from the UK Economic and Social Research Council through the Centre for the Microeconomic Analysis of Public Policy (CPP) at IFS (ES/M010147/1) and 'The impact of medical labour on variation in patient outcomes: evidence from English public hospitals' research grant (ES/S003118/1).

* Corresponding author.
E-mail addresses: george_s@ifs.org.uk (G. Stoye), max_w@ifs.org.uk (M. Warner).

[1] This includes many EU countries, such as Belgium, the Czech Republic, Estonia, Greece, Hungary, Ireland, Italy, Malta, Romania and Slovenia.

https://doi.org/10.1016/j.jebo.2023.06.011
0167-2681/

In this paper, we study the impact of a nationwide hospital strike that took place in England in 2016, and temporarily reduced the availability of 'junior' doctors in all public hospitals. Junior doctors – equivalent to Interns, Residents and Fellows in the US – comprise half of all hospital doctors in England and "play a vital role in delivering care to patients" (Health Education England, 2017). Working under the supervision of more senior doctors, they have the most regular contact with patients, monitoring and treating them on an hour-by-hour basis. We use detailed payroll and administrative patient data for the entire English public hospital system to quantify the impact of being exposed to this strike on hospital activity and patient outcomes, and to identify the patient groups most affected by the disruption.

Previous studies of doctor strikes generally conclude that they did not lead to worse patient outcomes (Cunningham et al., 2008; Slater and Ever-Hadani, 1983; Salazar et al., 2001; Erceg et al., 2007; Ruiz et al., 2013; Furnivall et al., 2018), or in some cases even improved outcomes (James, 1979; Roemer and Schwartz, 1979; Siegel-Itzkovich, 2000). This is in contrast to the increase in mortality rates during nursing strikes (Gruber and Kleiner, 2012), and recent evidence from Portugal of increased mortality and readmission rates during strikes by doctors and other healthcare workers (Costa, 2022).[2] However, a key challenge in estimating the impact of the strike is in controlling for potential changes in the unobserved severity of hospital patients during periods of disruption: for example, if sicker patients are temporarily diverted away from hospitals during strikes then average outcomes will improve even if care quality is unchanged. Comparisons of changes to patient outcomes before, during and after the strike, or between hospitals where workers have chosen to strike or not, may therefore lead to misleading conclusions about the impacts of the strikes on patients.

We contribute to this literature on the effects of doctor strikes by directly addressing this patient selection issue. Consistent with Furnivall et al. (2018) we find reductions in patient volumes on strike days, driven by both changes in the supply of and demand for healthcare. We also find significant changes in patient volumes on days around strikes, and in the observable characteristics of patients admitted to hospitals on strike days. These changes may bias any estimates of the effects of the strikes on patient outcomes using time-series variation alone.

Instead, we exploit additional cross-sectional variation in the exposure of different hospitals to the strike based on their pre-existing staffing mix to compare changes in patient outcomes on strike days across hospitals that were more or less affected by the strike. Furthermore, our rich patient data allows us to estimate the impact of doctor strikes on different patient groups to understand who is affected by short-term disruptions to medical care.

Specifically, we calculate the percentage of the doctor workforce in each hospital who were eligible to strike (i.e. the percentage of doctors who are 'junior'). Hospitals with a greater share of junior doctors would be expected to be more affected by the strike and may differentially cancel elective pre-planned activity. However, hospitals cannot turn away emergency patients and we would not expect patients to know the staffing structure of their local hospital or how affected it would be by the strikes. Any changes in emergency patient outcomes across hospitals should therefore be related only to how affected hospitals were by the strike itself rather than unobserved changes in patient severity. We validate this assumption by showing that there were no differences in the volumes or observable composition of emergency patients admitted to more or less exposed hospitals. We then estimate the impact of greater exposure to the strikes on patient outcomes, and how this impact varies across different patient groups.

In line with the majority of the existing literature, we find no statistically significant impacts of exposure to the strike on overall care quality: increased exposure to the strike did not lead to a significant increase in 30-day in-hospital mortality or readmission rates. Our estimates do indicate, however, that more exposed hospitals experienced much larger reductions in elective activity on strike days than less exposed hospitals. These results indicate that hospitals took strategic action to deal with strike disruption where possible, including (but not limited to) reducing non-urgent admissions. Taken together with the absence of large, adverse affects on average outcomes, this suggests that hospitals were generally able to mitigate the short-run staffing shock.

However, when carefully examining the impact of the strike across patient groups, we find evidence of negative impacts of the strike on some patient demographic groups. In particular, our estimates suggest that black patients experienced worse outcomes in more exposed hospitals on strike days compared with white patients, with these differences significant at either the 5% or 10% level depending on the specification used. Importantly, these differences occur *within* hospitals and conditional on a number of patient characteristics, and so are not explained by differential exposure across patient groups to striking doctors or by differences in clinical severity between groups. This suggests that greater exposure to the strikes had worse impacts for some specific minority groups, and is consistent with the worse outcomes found for minority groups in the US during periods when hospital capacity is strained (Singh and Venkataramani, 2022).

The primary contribution of this paper is to provide new evidence on the impact of doctor strikes on patient outcomes, addressing the selection issues in the current literature that we outlined above, examining heterogeneity in impacts across different patient groups, and extending the evidence base on strikes to a new universal healthcare system. This also builds on a small literature that examines the impact of wider health worker strikes (Gruber and Kleiner, 2012; Friedman et al., 2022; Costa, 2022).

More generally, our work relates to a literature that examines how patient outcomes change when hospital capacity is placed under strain (Evans and Kim, 2006; Hoe, 2022; Singh and Venkataramani, 2022). These studies tend to examine

[2] There is also emerging evidence of negative health consequences for patients during health worker strikes in developing countries. For example, Friedman et al. (2022) find that health worker strikes in Kenya - including doctors, nurses and wider healthcare staff - led to increased neonatal deaths.

shocks to hospital capacity due to patient demand as opposed to shocks to hospital supply. We extend this literature to examine how a shock to doctor labour supply impacted different patient groups within a national public healthcare system.

The rest of the paper is organised as follows. In Section 2, we set out the institutional setting and the background to the strikes. In Section 3, we describe our data. In Section 4, we show how patient numbers and composition changed during the strike, define our measure of strike exposure and provide evidence that differently exposed hospitals did not experience differential changes in their patient mix. In Section 5, we set out our empirical approach in detail, and in Section 6 show our results. Section 7 concludes.

2. Background

2.1. Hospital care in England

The vast majority of health care in England is publicly funded and free for all UK residents. Hospitals provide care to patients and are reimbursed by the government through a set of nationally agreed tariffs without patients paying any co-payments.

Unplanned or emergency hospital care is initially provided by emergency departments (EDs) which are part of large, publicly owned hospitals.[3] Patients arrive at these hospitals in two ways: either by ambulance, or as walk-in patients. Patients will usually attend their nearest hospital but walk-in patients can choose to attend any hospital if they wish so, and hospitals cannot turn away patients except in extremely rare circumstances. Patients are assessed upon arrival at the ED, and if necessary, can be admitted to the hospital as an emergency inpatient to receive treatment from specialists.[4] In addition to these large emergency departments, specialist emergency clinics or minor injury (‘walk in’) centres can treat patients with less serious diagnoses. There is no private market for emergency medicine, and so patients cannot substitute emergency care within the public sector to alternative providers.[5]

Planned or elective care is rationed through waiting times, and requires an initial referral from a primary care physician (known as a General Practitioner, or GP) or a senior hospital doctor (known as a consultant). A hospital consultant will usually assess the patient in an outpatient setting after this referral, with patients admitted at a later date for further treatment where required. Patients can choose which hospital they are initially referred to, and hospitals cannot refuse patients. However, hospitals will schedule admissions and can cancel pre-planned treatments if there is the lack of sufficient capacity to treat elective patients.

Hospitals are staffed by a mix of senior and junior doctors, alongside nurses and support staff. At the end of 2015, there were 96,338 doctors working in public general-acute hospitals, 49,701 (51.6%) of which were junior doctors and 46,637 (48.4%) were senior doctors. Appendix Table A1 shows the average characteristics of each staff group. This shows that junior doctors are on average younger, a higher proportion are women, and a similar proportion are white, compared to senior doctors.

The most senior doctors are known as consultants (equivalent to attending physicians in the US) and are legally responsible for their admitted patients. Consultants directly treat patients in addition to supervising more junior staff, and carrying out administrative tasks. A smaller group of senior doctors is Staff Grade, Associate Specialist and Specialty (SAS) doctors. SAS doctors are junior to consultants and their role tends to be more patient-focused. A small number of senior doctors are also employed on local or temporary contracts.[6] Together we refer to these doctors as senior doctors as they are no longer in training, and throughout the paper any counts of ‘senior’ doctors include both consultants and SAS doctors.

‘Junior’ doctors include all doctors who are in training, a heterogeneous group who vastly vary in experience and who are broadly equivalent to the group encompassing interns, residents and fellows in the US hospital system. This group includes Foundation Year doctors who work across a variety of different rotations in their first two years after medical school, and doctors in specialist training. Specialist training programmes vary in length, but can last up to eight years. As a result, some junior doctors have almost a decade of experience. Junior doctors have the most contact with patients, supervising and treating patients on an hour-by-hour basis.

2.2. The junior doctor contract dispute

The 2016 junior doctor strikes were the culmination of a long-running dispute between the government and medical unions about the introduction of a new junior doctor contract. The proposed new contract involved a number of changes, including higher basic pay but reduced pay for working at weekends and nights.[7] In September 2015 the British Medical Association (BMA), the primary trade union for doctors, announced it would ballot junior doctors in England for industrial

[3] EDs are known as Accident and Emergency (A&E) departments in England. Groups of hospitals that share management are grouped into NHS ‘trusts’. Throughout the paper we refer to trusts as ‘hospitals’.

[4] Emergency patients do not have pre-assigned hospital doctors, and will be treated by the staff working in the hospital at time of treatment.

[5] NHS primary care physicians do sometimes provide some basic treatments to emergency patients but this is not comparable to emergency hospital treatment.

[6] Doctors with unknown contracts are excluded from our analysis. Doctors with an ad-hoc local contract are included as senior doctors.

[7] https://www.bbc.co.uk/news/health-34775980

action. In November 2015 the BMA announced that it had balloted more than 37,000 junior doctors and that 98% voted in favour of full strike action.[8]

This resulted in five strikes that took place in all English public hospitals between January and April 2016. Each strike lasted 24 or 48 hours, and took place on a Tuesday, Wednesday or Thursday.[9] In the first four strikes, junior doctors continued to provide emergency care to hospital patients, with only pre-planned (elective) care directly affected. However, despite initially intending to repeat this on the final strike dates, the BMA announced in late March that it would include a "full withdrawal of labour".[10] This meant that emergency care was also withdrawn by participating junior doctors. In our analysis we therefore classify the first four strikes as elective strikes, and the final strike as an all-out strike. As we discuss in detail in Section 4.2, hospitals have little control over the number of emergency patients who seek treatment on any given day. This is in contrast to elective care, where treatments can be postponed by the hospital. For this reason, our primary results focus on the effects of the final strike on emergency patient outcomes only.

Junior doctors were not obliged to strike, although the majority of them did. In the first four strikes, NHS England reported that between 39% to 46% of junior doctors reported for duty on the day shift (NHS England, 2016a; 2016b; 2016c; 2016d). This includes many who were rostered for emergency services and so could not strike, and those who were absent for other reasons. Subsequent analysis by NHS England estimated that 88% of those who did not report for duty during the final elective strike in April were striking (NHS England, 2016e). During the two days of all-out strike, NHS England estimated that 78% of junior doctors who were expected to be working did not report for duty (NHS England, 2016e; 2016f).[11]

Hospitals responded to the strikes in a number of ways and a national response plan was established by NHS England to mitigate the negative consequences of the strikes. This included cancelling many elective operations and outpatient appointments on strike days to free up staff time, with over 100,000 outpatient appointment cancellations and more than 25,000 fewer planned admissions than expected as a result of the strikes (Furnivall et al., 2018). Senior doctors and nurses from within hospitals were redeployed to provide emergency cover, and freelance locum doctors were temporarily employed.[12] In addition, some hospitals cancelled holidays and study leave for other staff groups.[13] In advance of the all-out strike, NHS England also asked primary care providers and ambulance trusts to provide additional support to reduce the demand for hospitals.[14] Our estimates of the impact of the strikes will therefore include the impact of hospital responses aimed at mitigating the effects of the strikes.[15]

Some of these responses may be common to hospitals or healthcare systems around the world when facing strikes, such as cancelling or postponing non-urgent activity in order to reduce the burden on staff working during the strike and reallocate available resources to substitute for striking staff. But other responses, such as using resources from primary care and ambulance providers, relied on the integration of the English health system. That would suggest that our analysis provides a lower bound on the potential effects of a *national* strike of equivalent doctors in other healthcare systems. Hospital responses to strikes also depend on the amount of notice given. Strikes that are announced with less notice than the strikes we examine are likely to have larger impacts on patient outcomes if hospitals cannot take mitigating actions in time. But if the strikes only took place among certain employers in another healthcare system, it may be that patients or staff could be temporarily redistributed between providers which could further mitigate the impacts of doctor strikes.

3. Data

3.1. Payroll data

Our source of data on hospital staffing is the Electronic Staff Record (ESR). These data are the monthly payroll for all staff directly employed by the NHS, which we use for December 2015 only, the month before the first strike. These data include an occupational code, demographic characteristics, a breakdown of monthly pay (e.g basic pay, overtime, geographic allowances etc.) and the number of hours or shifts worked for each employee.

We identify doctors and their grades using staff groups and national pay codes in the data.[16] Doctors are only classed as working for a given hospital in a period if they earn strictly positive pay. As doctors may work for multiple hospitals, we assign each doctor a primary hospital, based on the hospital where they have the highest pay.

[8] https://www.theguardian.com/society/2015/nov/19/nhs-strikes-junior-doctors-vote-action-bma

[9] These strikes took place on January 12th (Tuesday), February 10th (Wednesday), March 9-10th (Wednesday-Thursday), April 6-7th (Wednesday-Thursday) and April 26-27th 2016 (Tuesday-Wednesday).

[10] https://www.theguardian.com/society/2016/mar/23/junior-doctors-to-escalate-industrial-action-to-all-out-strike-next-month

[11] Data on how many junior doctors went on strike at each hospital are unavailable.

[12] https://www.theguardian.com/society/2016/apr/25/consultants-locums-called-into-a-and-e-for-junior-doctors-strike

[13] https://www.bbc.co.uk/news/health-36134103

[14] https://www.england.nhs.uk/2016/04/industrial-action-update/

[15] Many of the potential hospital responses are not recorded in any publicly available data. Hospitals do not publish (and often do not collect) day-to-day information on the use of temporary staff, staff absences or non-paid overtime. We therefore cannot directly model variation in the response to the strikes by different hospitals, or estimate the impact of these actions on patient outcomes.

[16] Doctors are classified as junior doctors if they have any of the the pay codes for foundation doctors or registrars.

Table 1
Hospital summary statistics.

	Mean	S.D.
Elective admissions per day	181.5	160.3
Emergency admissions per day	115.1	56.5
Junior doctors	394.5	419.2
Senior doctors	370.1	206.5
Nurses	2,009.4	1,098.9
Number of hospitals	126	

Notes: Staff numbers are for December 2015. Average admissions are for our whole sample period.

Table 2
Patient summary statistics.

	Elective		Emergency	
	Mean	S.D.	Mean	S.D.
Age	57.3	20.5	52.6	28.5
% Female	51.3%	50.0%	52.3%	49.9%
Charlson	1.32	2.04	1.14	1.81
Days in hospital in last year	5.13	19.8	7.77	22.4
30-day readmissions	0.0626	0.242	0.152	0.359
30-day in-hospital mortality	0.00297	0.0544	0.0373	0.190
Number of patients	50,117,377		31,786,819	

3.2. Hospital records

Our source of data on patient volumes, characteristics and outcomes is the Hospital Episode Statistics (HES). These contain the administrative patient-level hospital records for all care received in public hospitals in England. Each record contains information on dates and modes of admission (elective or emergency) and discharge, hospital, clinical speciality and consultant identifiers, a set of patient characteristics and up to 20 diagnoses codes.

We define two measures of patient outcomes. We create an indicator of an emergency readmission if a patient is observed having an emergency admission within 30 days of the original admission. We also create an indicator of in-hospital mortality if the patient is recorded as dying in any NHS hospital within 30 days of an admission.

3.3. Sample selection

Our sample period is 1st April 2012 to 31st March 2017. We restrict our sample to general acute NHS hospitals. This excludes private hospitals providing care for NHS patients, non-acute hospitals (such as community and mental health hospitals) and specialist acute hospitals. At the patient level, we also remove privately-funded patients in NHS hospitals from our sample and those with missing demographic information. This provides an analysis sample of 50.1 million elective and 31.8 million emergency admissions across 126 hospitals.

Table 1 shows summary statistics for our sample of 126 hospitals. The mean hospital had 182 elective admissions and 115 emergency admissions per day during our sample period. In December 2015, the month before the strikes started, the mean hospital had 395 junior doctors and 370 senior doctors. As a result, an average of 51.6% of employed doctors were junior doctors.

Table 2 shows summary statistics for our sample of elective and emergency patients. The mean 30-day readmissions rate is 6.3% for elective patients and 15.2% for emergency patients. The mean 30-day in-hospital mortality rate is 3.7% for emergency patients and 0.3% for elective patients. Throughout our analysis we use Charlson scores as a measure of patient severity (Quan et al., 2005). Charlson scores measure the number and severity of comorbidities that each patient has recorded.

4. Impact of the strikes at hospital level

4.1. Aggregate evidence

We begin by analysing the impacts of the strikes on the daily volume of elective and emergency admissions at public hospitals. Figure 1 shows the national volume of elective and emergency patients admitted on each day between 1st January 2016 and 30th April 2016. The vertical lines represent the strike periods. It shows that on all-strike days, there were substantial falls in elective admissions: on average, the total number of elective admissions was 4.7% lower on an elective strike day and 13.2% lower on an all-out strike day than on all other weekdays in the same period. There was also a 7.3% fall

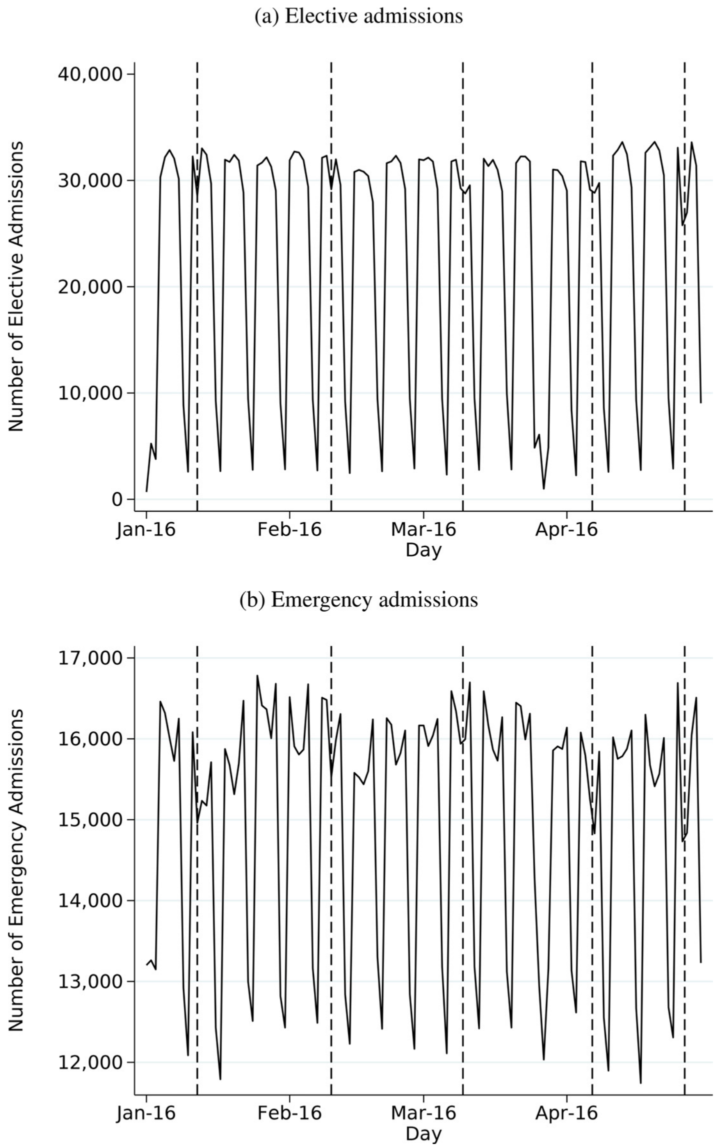

Fig. 1. National number of admissions.

in emergency patients on emergency strike days and a smaller 3.4% fall on elective strike days. Noticeably, these falls are much smaller than the drop in admissions at weekends, but represent substantial deviations from usual weekday patient volumes. Figure 1 shows that there is a large amount of day-to-day fluctuations in the number of admissions. To adjust for these fluctuations, we estimate the following specification at the hospital level:

$$Y_{jrt} = \boldsymbol{\alpha}_j + \boldsymbol{\psi}_{f(t)} + \delta_1 ElStrike_t + \delta_2 AllStrike_t + \boldsymbol{\rho}(region_r \times year_t) + \boldsymbol{\sigma}(type_j \times year_t) + \epsilon_{it} \tag{1}$$

where Y_{jrt} is the log number of emergency department arrivals, elective or emergency admissions at hospital j in region r on day t. $\boldsymbol{\alpha}_j$ is a hospital fixed effect, and controls for permanent differences in admissions across hospitals. The inclusion of $\boldsymbol{\psi}_{f(t)}$ allows us to flexibly control for seasonality in admissions, and includes month by year and day of week by year fixed effects, as well as fixed effects for a set of holidays. $ElStrike_t$ is a dummy variable that takes the value of one for days when a strike affected elective but not emergency activity, and zero otherwise. $AllStrike_t$ is a dummy variable that takes the value of one when a strike affected both elective and emergency activity. We also include regional $(region_r \times year_t)$ and hospital type $(type_t \times year_t)$ time trends to control for broader time trends in admissions across regions and hospital type respectively. Standard errors are clustered at the hospital level.

Table 3 shows the results.[17] Column one reports results for emergency department arrivals, and shows that they fell by 5.0% on elective strike days and 13.3% on all-out strike days. Column two repeats this analysis for emergency admissions, and shows a small fall of 2.5% on days with an elective strike and a larger fall of 6.5% on days when emergency services

Table 3
Log volumes.

	(1) Emergency Arrivals	(2) Emergency Admissions	(3) Elective Admissions
Elective Strike	−0.0499***	−0.0249***	−0.104***
	(0.00223)	(0.00364)	(0.00823)
All Strike	−0.133***	−0.0648***	−0.203***
	(0.00547)	(0.00692)	(0.0182)
N	273,854	276,066	271,038

Note: (1) Standard errors clustered at the hospital level are displayed in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. (2) The outcome variables in column 1 is the log number of patients arriving at emergency departments per hospital per day. (3) The outcome variable in column 2 is the log number of patients admitted via emergency admission methods per hospital per day. (4) The outcome variable in column 3 is the log number of patients admitted via elective admission methods per hospital per day. (5) A small number of hospital-day observations have zero admissions or arrivals and we drop these from the analysis, see footnote [17] for more details. (6) All regressions include hospital fixed effects, a set of time fixed effects and regional and hospital-type trends.

Table 4
Log volumes around strike days.

	(1) Emergency Arrivals	(2) Emergency Admissions	(3) Elective Admissions
Day Before Elective Strike	−0.0168***	0.00571	0.0191***
	(0.00287)	(0.00397)	(0.00674)
Day After Elective Strike	−0.00578**	0.00128	0.0141
	(0.00269)	(0.00422)	(0.00882)
Day Before All Strike	0.00644	0.0295***	0.0212*
	(0.00537)	(0.00853)	(0.0121)
Day After All Strike	−0.0173***	0.0236***	0.0642***
	(0.00620)	(0.00819)	(0.0104)
N	273,605	275,814	270,808

Note: (1) Standard errors clustered at the hospital level are displayed in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. (2) The outcome variables in column 1 is the log number of patients arriving at emergency departments per hospital per day. (3) The outcome variable in column 2 is the log number of patients admitted via emergency admission methods per hospital per day. (4) The outcome variable in column 3 is the log number of patients admitted via elective admission methods per hospital per day. (5) A small number of hospital-day observations have zero admissions or arrivals and we drop these from the analysis. (6) All regressions include hospital fixed effects, a set of time fixed effects and regional and hospital-type trends.

were directly affected.[18] Column three reports results for elective admissions, and shows that elective admissions fell on days with an elective strike by 10.4%. The fall in elective admissions was greater, at 20.3%, on days when the strike also covered emergency activity.

Our results are similar to the results of Furnivall et al. (2018), who found a 14.7% reduction in emergency department arrivals, a 7.8% reduction in emergency admissions and a 19.9% reduction in elective admissions on the all-out strike days. Their analysis used the same data but a different methodology, comparing activity on strike days with the average weekday activity on weeks around the strikes.

Table 4 repeats this analysis with additional dummy variables for the days before and after each type of strike. Elective admissions were 2% higher on the days before the strikes, and 6% higher on the day after the all-out strike. This is consistent with elective activity being displaced by the strikes and hospitals rescheduling care. There was no significant change on days after the elective strikes. Emergency admissions were unchanged around the elective strikes but were 3% and 2% higher the day before and day after the all-out strike, respectively. This suggests that some of the 'missing' emergency admissions during the all-out strikes took place in advance or after the strikes. Emergency department arrivals were 2% lower the days before the elective strikes, and 0.6% lower the days after. They were also 2% lower the day after the all-out strike, and unchanged the day before. This suggests that overall, emergency department arrivals were reduced during the whole strike period.

These results indicate that hospitals services were significantly disrupted by the strikes, with large falls in the number of patients admitted on these days, and smaller changes around the days of strikes. Patient care-seeking behaviour also changed, with fewer arrivals at emergency departments on strike days. Our ultimate objective is to study how the care

[18] Emergency department arrivals are the number of people who arrive and seek care at emergency departments. Many will receive treatment within the department and will not be admitted as emergency patients into the hospital. There are other routes for emergency admissions, such as through primary care doctors or other hospital doctors, but these account for only a small fraction of emergency admissions.

Table 5
Elective patient composition.

	(1) Age	(2) Female	(3) Depr.	(4) Usage	(5) Charlson	(6) Pred. Read.	(7) Dist.
Elective Strike	0.410***	−0.00114	−0.0484	0.142**	0.0764***	0.00113***	14.76
	(0.112)	(0.00196)	(0.0453)	(0.0704)	(0.0102)	(0.000351)	(69.54)
All Strike	1.262***	0.00590	−0.0786	0.659***	0.148***	0.00199***	−98.02
	(0.220)	(0.00382)	(0.0813)	(0.212)	(0.0235)	(0.000661)	(120.0)
N	271,038	271,038	270,960	271,038	271,038	271,038	270,995

Note: (1) Standard errors clustered at the hospital level are displayed in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. (2) The outcome variables in column 1 is the mean age of patients admitted by elective admission methods per hospital trust per day. (3) In column 2 the outcome variable is the percentage of elective admissions that were female. (4) In column 3 the outcome variable is the average index of multiple deprivation (IMD) score of the local areas (LSOA) where admitted patients live. (5) In column 4 the outcome variable is the mean number of days that admitted patients have previously spent in hospital over the previous year. (6) In column 5 the outcome variable is the mean Charlson score for admitted patients. (7) In column 6, the outcomes variable is the mean predicted readmissions rate. (8) In column 7 the outcome variable is the average distance between the centroid of the patient's home LSOA and the centroid of the hospital's primary site measured in metres. (9) All regressions include hospital fixed effects, a set of time fixed effects and regional and hospital-type trends.

Table 6
Emergency patient composition.

	(1) Age	(2) Female	(3) Depr.	(4) Usage	(5) Charlson	(6) Pred. Read.	(7) Dist.
Elective Strike	0.357***	0.00411**	0.0878*	0.125	0.0133*	0.000338	118.1
	(0.111)	(0.00164)	(0.0462)	(0.0939)	(0.00774)	(0.000210)	(79.81)
All Strike	−0.991***	0.00111	0.0355	0.145	−0.0373***	−0.000239	−225.2
	(0.214)	(0.00300)	(0.0957)	(0.156)	(0.0119)	(0.000367)	(149.4)
N	276,066	276,066	276,055	276,066	276,066	276,066	276,055

Note: (1) Standard errors clustered at the hospital level are displayed in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. (2) See notes for Table 5; definitions apply for patients admitted by emergency admission method.

quality of patients who were still admitted was affected by the strikes.[19] However, these large changes in the volume of patients raises the issue of whether patient composition was also affected. For example, we would expect that patients with less severe conditions would be most affected as their care is likely to be less urgent, and so observed patient outcomes would be worse on average even if care quality was unaffected. And while changes in observed patient casemix can be controlled for directly in our analysis, any large changes in such characteristics would raise concerns about other unobserved changes in severity that could bias our estimates of the impacts of the strikes on care quality.

To examine whether observed patient casemix changed as a result of the strikes, we re-estimate Eq. (1), replacing volumes on the left hand side with the mean characteristics of admitted patients in hospital j in region r on day t. Table 5 shows the results of this exercise for elective patients. In column one, the dependent variable is mean age. Columns two to six show results for the percentage of patients who are female, the mean deprivation score of the patient's local area of residence, the mean number of days spent by a patient in hospital in the past year, the mean Charlson score, and the mean 30-day predicted readmissions rate.[20] The final column shows results for the average distance from each patient's home address to the primary address of each hospital. This is calculated as the distance (in metres) between the centroid of the Lower Super Output Area (a small geographical census unit) of the patient's home address and the equivalent centroid of the hospital's primary address.[21] Table 6 repeats this for emergency patients.

On both elective and all-out strike days elective patients were on average older, had higher prior hospital usage, had a higher average Charlson score and a higher predicted readmission rate. All of these factors suggest that elective patients were sicker on average during strike days. Ex-ante it is not obvious how we would expect elective patient severity to change during the strike: we might expect average sickness to decrease if hospitals decided to delay more risky or complex surgery

[19] Another group affected by the strike are those who would have been admitted if there had been no strike but were not admitted because of the strike. The fall in admission volumes on strike days shows that this is a sizeable group. However, we cannot identify this set of individuals from our data and so cannot estimate the effect of the strikes on their outcomes.

[20] Deprivation is measured by the 2011 Index of Multiple Deprivation Score at the patient's Lower Layer Super Output Area (a small administrative area with an average of 1600 people in 2011). Predicted readmission rate is calculated by regressing a dummy variable for 30-day readmission on age by gender dummies, primary diagnosis dummies and Charlson dummies, excluding strike days.

[21] We exclude patients whose home address is not in England, or who have no recorded home address. This is the case for 306,000 (0.6%) elective patients and 361,000 (1.1%) emergency patients in our sample.

on strike days. This result therefore reveals suggests that hospitals prioritised sicker elective patients during the strike rather than delaying their care to a later date.

The pattern for emergency patients is less clear. On elective strike days, emergency patients were older on average, more likely to be female, more likely to come from deprived areas and had a higher Charlson score. On all-out strike days, the average age and Charlson score were significantly lower, which could suggest that patient severity may have been lower on strike days.[22] It is important to remember, however, that this is the composition of those who were actually admitted to hospital, rather than those who arrived to emergency departments.[23] The larger reduction in emergency department arrivals than admissions could suggest that those arriving at emergency departments were sicker on average. But it could also suggest that the reduction in arrivals was larger than the reduction in capacity and so hospitals actually admitted less severe patients than they normally would have, consistent with the reduction in some of our severity measures.

For both elective and emergency admissions, there is no change in the average distance between the patient's home address and the primary address of the hospital they are admitted to. This suggests there was no systematic change in the hospitals that patients went to during the strikes or that ambulances systematically diverted patients to other hospitals.

Taken together, these results suggest that the strikes led to large changes in the type of both elective and emergency patients treated on strike days. As a result, we cannot rely solely on time variation to study the impacts on admitted patients due to concerns about unobserved changes in patient severity. We therefore now consider an additional source of cross-sectional variation in exposure to the strikes that is plausibly unrelated to changes in patient mix of emergency patients in order to address this issue.

4.2. Variation by pre-existing staff mix

4.2.1. Defining strike exposure

Throughout the rest of the paper, we exploit differences in the composition of the pre-existing workforce to examine the impact of being exposed to the strikes on emergency patient outcomes. Intuitively, hospitals should be affected differently by the strikes depending on how easily they can substitute junior doctors with other staff (such as senior doctors). For example, hospitals with a greater share of junior doctors among their total doctor workforce would expect to be more affected by the strike as a larger share of their workforce is eligible to strike.

Hospitals are likely to know their relative exposure to the strike, and react accordingly to reduce pressure. One margin that hospitals can control is pre-planned elective activity. We would therefore expect that hospitals with a greater share of junior doctors (and who are therefore more exposed under our measure) would cancel more elective procedures in order to free up capacity for the treatment of emergency patients during this period. We therefore cannot exploit the differences in pre-strike staffing mix to determine the impact on elective patients, but we can use this to see whether hospitals we expect to be more affected by the strike do take more mitigating action.

However, while national efforts were made to reduce total demand for emergency care on strike days (as reflected by the drop in national emergency attendances and admissions in Table 3), individual hospitals were unlikely to be able to respond by changing emergency admissions as they cannot choose to turn away emergency patients. Moreover, patients are unlikely to be aware of the staffing situation at their local hospital and so would not take this into account when choosing which hospital to seek treatment from if required during the strike.[24] Any changes in emergency patient outcomes across hospitals should therefore be related only to how affected hospitals were by the temporary loss of staff rather than unobserved changes in patient severity. We test the plausibility of this assumption in the next section.

Specifically, we define the following strike exposure measure for each hospital j:

$$Exp_j = \frac{Junior_j}{Junior_j + Senior_j} \tag{2}$$

where $Junior_j$ and $Senior_j$ are the number of junior and senior doctors employed by hospital j in December 2015. This provides a measure of the fraction of the doctor workforce employed at the end of the previous year who were eligible to strike.

Figure 2 shows how exposure varies across hospitals. The median value is 0.47, an almost equal mix of senior to junior doctors. However, this varies considerably across hospitals, with an exposure of 0.58 at the 90th percentile compared to an exposure 0.29 at the 10th percentile. A number of institutional and historical factors determine the relative employment of junior doctors to senior doctors. Teaching hospitals (hospitals associated with university medical departments) employ a higher proportion of junior doctors because these hospitals are more involved in the training and education of doctors.

[22] Charlson scores measure pre-existing co-morbidities but this may not always be a relevant measure of severity for some types of emergency admissions, such as accidents.

[23] The emergency department data has much less detail on patient diagnoses than the inpatient records, making it hard to assess how severity changed.

[24] Two other concerns are that ambulance services may re-distribute patients between hospitals and that patients may substitute between elective and emergency care. In the next section we show the strike did not lead to changes in the average distance travelled by patients for care, suggesting that patients were not redistributed across hospitals. The second concern seems less plausible given the temporary nature of the strikes and the institutional set-up of the NHS. Patients waiting for elective care will generally already have waited weeks for treatment, suggesting that while treatment is required it is not required on a specific day. In addition, patients would not be able to access alternative treatment at emergency departments (e.g. if patients were waiting for elective surgery, they would not be able to get such care at the ED on the day of the strike).

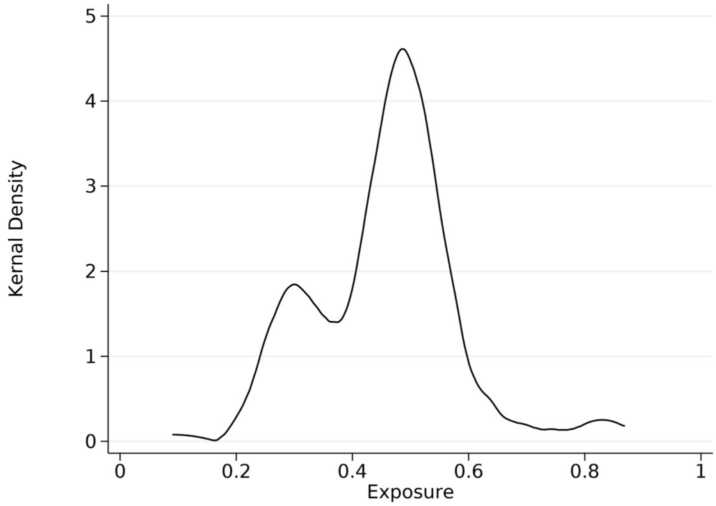

Fig. 2. Distribution of exposure. *Note:* Exposure is as defined in Eq. (2) using data from December 2015.

Table 7
Characteristics by hospital exposure.

	Below Median Exposure	Above Median Exposure	Difference
Elective admissions per day	157.6	205.5	−47.9 (19.7) **
Emergency admissions per day	102.0	128.3	−26.3 (9.08) ***
Junior doctors	187.6	601.3	−413.7 (65.1) ***
Senior doctors	307.6	432.7	−125.1 (35.2) ***
Average age of doctors	41.1	38.7	2.42 (0.247) ***
% Female doctors	39.1%	46.7%	7.58 (1.03) ***
% White doctors	57.2%	58.1%	0.902 (2.45)
Age	53.0	52.0	0.718 (0.685)
Charlson	1.14	1.14	−0.000505 (0.227)
Days in hospital last year	7.52	7.97	−0.452 (0.178) **
30-day readmissions	0.149	0.155	−0.00529 (0.00180) ***
30-day in-hospital mortality	0.0393	0.0358	0.00350 (0.00108) ***

Note: (1) All outcomes are for our whole sample period, apart from numbers and characteristics of doctors which are for December 2015 only. (2) Standard errors for differences clustered at the hospital level are displayed in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Hospitals in more urban areas employ a greater proportion of junior doctors than hospitals in more rural areas, perhaps because younger doctors prefer to live in more urban areas. Larger hospitals also employ a greater proportion of junior doctors. But even within similar areas, there is a large degree of heterogeneity in the proportion of junior doctors employed by nearby hospitals. In London, for example, the proportion of junior doctors ranges from 45% to 69%. Unobserved hospital-level institutional and historical factors are therefore also important determinants of the staffing mix.

Table 7 shows how hospital characteristics and patient characteristics varied between hospitals with above and below-median exposure. Less exposed hospitals tend to be smaller, both in terms of their patient volumes and staffing levels. Consistent with the characteristics of junior and senior doctors reported in Table A1, more exposed hospitals have a younger and more female doctor workforce. There is no significant difference in doctor ethnicity by exposure.

At the patient level, there is no significant difference in the average age or Charlson score of patients, although patients admitted to more exposed hospitals had higher prior hospital usage. Patients admitted to hospitals with below median exposure have lower 30-day readmission rates and higher 30-day in-hospital mortality rates on average. These differences highlight the need for our empirical strategy to control for any persistent differences in outcomes across hospitals that were differently exposed to the strikes.

4.2.2. Examining patient composition by strike exposure

Using this exposure measure, we examine whether there was a differential impact of the strikes on patient volumes and average patient characteristics across differently exposed hospitals by estimating the following specification:

$$\begin{aligned} Y_{jrt} = {} & \boldsymbol{\mu}_j + \boldsymbol{\pi}_{f(t)} + \gamma_1 ElStrike_t + \gamma_2 AllStrike_t + \gamma_3 (ElStrike_t \times Exp_j) \\ & + \gamma_4 (AllStrike_t \times Exp_j) + \boldsymbol{\phi}(region_r \times year_t) + \boldsymbol{\chi}(type_j \times year_t) + \upsilon_{it} \end{aligned} \quad (3)$$

Table 8
Log volumes.

	(1) Emergency Arrivals	(2) Emergency Admissions	(3) Elective Admissions
Elective Strike	−0.105***	−0.0199	0.0272
	(0.0310)	(0.0233)	(0.0498)
Elective Strike * Exposure	0.120*	−0.0108	−0.287***
	(0.0674)	(0.0505)	(0.107)
All Strike	−0.172***	−0.0725**	0.107
	(0.0380)	(0.0313)	(0.0838)
All Strike * Exposure	0.0847	0.0167	−0.678***
	(0.0819)	(0.0673)	(0.177)
N	273,854	276,066	271,038

Note: (1) Standard errors clustered at the hospital level are displayed in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. (2) The outcome variables in column 1 is the log number of patients arriving at emergency departments per hospital per day. (3) The outcome variable in column 2 is the log number of patients admitted via emergency admission methods per hospital per day. (4) The outcome variable in column 3 is the log number of patients admitted via elective admission methods per hospital per day. (5) A small number of hospital-day observations have zero admissions or arrivals and we drop these from the analysis, see footnote [25] for more details. (6) All regressions include hospital fixed effects, a set of time fixed effects and regional and hospital-type trends.

where Y_{jrt} is either the (log) number, or an average characteristic, of patients admitted to hospital j in region r on day t. The specification is identical to Eq. (1), with the addition of two interaction terms between the strike dummies and the exposure measure. In this case, our coefficient of interest is γ_4: if our approach is valid we would expect this coefficient not to be significantly different from zero for emergency patients. Standard errors are again clustered at the hospital level.

Table 8 shows the results for patient volumes.[25] In column one, the dependent variable is log emergency department arrivals. The results show that volumes fell on strike days, and that there was no statistically significant differential effect by exposure on the all-out strike days.[26] In column two, the dependent variable is log emergency admissions. The results show that there was no impact of elective strike days on emergency volumes, and there was no additional effect depending on how exposed hospitals were to the strike. For emergency strike days, there was a 7.3% reduction on emergency admissions, but there was no statistically significant additional impact of being more exposed to the strike. In both cases, the estimated coefficients are close to zero, and are much smaller than the point estimate for emergency arrivals in column one. In line with our hypothesis, this suggests there was no differential changes in the numbers of emergency patients admitted to hospitals with a greater or lesser share of junior doctors in their pre-strike doctor workforce as a result of the strike.

Column three repeats this analysis for elective admissions. In contrast to the results in Table 3, there is no longer a reduction in elective activity associated with an elective or an emergency strike. However, there is now a large and statistically significant negative coefficient on both the interaction terms. This suggests that the reductions in elective activity were driven by hospitals with a greater share of junior doctors. On the emergency strike days, our results suggest that hospitals with a 10 percentage point higher share of junior doctors cancelled 6.8% more elective admissions. This is consistent with the ability of hospitals to temporarily postpone elective procedures, with hospitals more affected by the strikes more likely to cancel such operations. This also provides reassuring evidence that our measure of exposure does capture hospitals that were put under greater pressure from the strikes as a result of their overall available staff resources. For this reason we only estimate the impact of the strikes on emergency patients.

Appendix Table B2 repeats this analysis for the days immediately before and after the strikes. This shows no differential changes in emergency admissions either the day before or after the all out strikes. The table shows that more exposed hospitals had slightly higher emergency department arrivals the day after the strike: a 10 percentage point increase in exposure is associated with a 1.5% increase in arrivals. Since there is no reduction in arrivals by exposure the day before or the day of strikes this likely does not signify strategic changes in arrivals. Instead, it may represent a consequence of the greater disruption at these hospitals during the strikes. Hospitals that were more exposed also had significantly lower elective admissions the day after the all-out strikes, suggesting the cancellation of elective activity by hospitals continued the day after the strikes.

Table 9 shows results for a number of characteristics of emergency patients. In all cases, there are no statistically significant differences in the changes in the composition of emergency patients on all-out strike days across differently exposed hospitals. There is also no change in the average distance travelled by patients for treatment. This suggests that patients did

[26] The point estimate suggests that a 10% increase in exposure led to a 0.8% increase in arrivals on a strike day but is quite imprecisely estimated. There is some evidence of increased ED attendances on elective strike days for more exposed hospitals, significant at the 10% level. We also repeat the analysis separately for walk-in ED patients and ambulance arrivals. We find no statistically significant impacts of exposure to the all-out strike on either patient types, suggesting that more exposed hospitals did not divert ambulance patients to other less-exposed hospitals on strike days. Full results are available upon request.

Table 9
Emergency patient composition.

	(1) Age	(2) Female	(3) Depr.	(4) Usage	(5) Charlson	(6) Pred. Read.	(7) Dist.
All Strike	−0.568	0.00234	0.405	−0.189	−0.0272	−0.000618	−270.0
	(0.739)	(0.0101)	(0.424)	(0.667)	(0.0441)	(0.00146)	(374.5)
All Strike * Exposure	−0.925	−0.00269	−0.807	0.731	−0.0221	0.000829	98.09
	(1.531)	(0.0205)	(0.816)	(1.439)	(0.0889)	(0.00298)	(760.8)
N	276,066	276,066	276,055	276,066	276,066	276,066	276,055

Note: (1) Standard errors clustered at the hospital level are displayed in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. (2) The outcome variables in column 1 is the mean age of patients admitted by emergency admission methods per hospital trust per day. (3) In column 2 the outcome variable is the percentage of emergency admissions that were female. (4) In column 3 the outcome variable is the average index of multiple deprivation (IMD) score of the local areas (LSOA) where admitted patients live. (5) In column 4 the outcome variable is the mean number of days that admitted patients have previously spent in hospital over the previous year. (6) In column 5 the outcome variable is the mean Charlson score for admitted patients. (7) In column 6 the outcome variable is the mean predicted readmissions rate. (8) In column 7 the outcome variable is the average distance between the centroid of the patient's home LSOA and the centroid of the hospital's primary site measured in metres. (9) All regressions include hospital fixed effects, a set of time fixed effects and regional and hospital-type trends.

not systematically go to different hospitals - either by their own choice or because ambulances diverted them - in a way that is correlated with our measure of exposure to the strikes.

Taken together, these results suggest that there was no differential impact on the volume of emergency patients - or the observed characteristics of these admitted patients - across hospitals that were plausibly more or less exposed to the strike due to their pre-strike staffing mix. We therefore now set out an empirical approach to examine the impact of strike exposure on the short-run outcomes of emergency patients admitted to hospital during the strikes.

5. Empirical strategy

To examine the impact of being exposed to the strikes on the health outcomes of emergency patients, and to exploit the richness of the health information available at the patient level in the health records, our main specification is at the patient level. Specifically, we estimate the following baseline specification:

$$\begin{aligned} Y_{ijrt} = {} & \boldsymbol{\zeta}_j + \boldsymbol{\lambda}_{f(t)} + \beta_1 ElStrike_t + \beta_2 AllStrike_t + \beta_3 (ElStrike_t \times Exp_j) \\ & + \beta_4 (AllStrike_t \times Exp_j) + \boldsymbol{\theta} \boldsymbol{X}_{ijt} + \boldsymbol{\eta}(region_r \times year_t) + \boldsymbol{\tau}(type_j \times year_t) + \nu_{it} \end{aligned} \tag{4}$$

where Y_{ijrt} is the outcome for an emergency patient i who is admitted to hospital j in region r on day t. We again control for hospital fixed effects ($\boldsymbol{\zeta}_j$), time effects ($\boldsymbol{\lambda}_{f(t)}$) and regional and hospital-type trends ($region_r \times year_t, type_j \times year_t$) in the same way as in the hospital-level Eq. (3). In addition, since the unit of analysis is now at the patient level, we include a vector of detailed controls for the observable health of each patient ($\boldsymbol{X}_{ijt}$), including age by gender variables, Charlson dummy variables and primary diagnosis fixed effects. Standard errors are clustered at the hospital level.

Coefficients β_1 and β_2 capture how patient outcomes changed on strike days in hospitals if they had zero exposure. This can be interpreted as the national change in patient outcomes on strike days that is common to all hospitals. Coefficients β_3 and β_4 capture the effect of exposure on elective and all-out strike days respectively. In our results we focus on the impact of exposure to the all-out strike (β_4) as the all-out strike alone directly affected emergency patients.[27] A positive coefficient, for example, would imply that more exposed hospitals had worse patient outcomes than less exposed hospitals. Our identifying assumption is that, conditional on the controls, the interaction between our exposure measures and the strike dummies are uncorrelated with the error term, ν_{it}.

This identifying assumption would be violated if there are unobserved factors that affect patient outcomes which are also correlated with the interaction terms. For example, if there were differential trends in hospital activity over time correlated with exposure, this would bias the estimated impact of strikes. Two likely sources of differential trends come from regional and hospital type differences. For example, teaching hospitals are likely to have higher exposure and may have differential trends in activity compared to non-teaching hospitals. For this reason we include region and teaching hospital year trends in our main specification. In Section 6.3 we also repeat our analysis using a much shorter time period around the strikes, similar to the methodology of Furnivall et al. (2018), and find similar results to our main specifications.

It is important to note that our approach only allows us to identify whether more or less exposed hospitals were affected differently by the strike as opposed to identifying the *total* effect of the strike on patients. Attempting to estimate the overall effect of the strike would require us to make the more restrictive identifying assumption that the error term is uncorrelated with the strike dummy variables, or in other words, that patient composition was unchanged on strike days

[27] We include the terms relating to elective strikes to control for potential differences in outcomes caused by these earlier strikes but do not report the coefficients for the sake of brevity. Results are available upon request.

Table 10
Patient outcomes.

	(1) Readmissions	(2) Mortality
All Strike	−0.000954	−0.000560
	(0.00790)	(0.00434)
All Strike * Exposure	0.0136	0.00279
	(0.0154)	(0.00868)
Outcome Mean	0.152	0.0373
N	31,786,339	31,786,339

Note: (1) Standard errors clustered at the hospital level are displayed in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. (2) The outcome variable in the first column is whether the patient has an emergency admission within 30 days of their original admission. (3) The outcome variable in the second column is whether the patient dies in hospital within 30 days of their original admission. (4) All regressions include hospital fixed effects; a set of time fixed effects; regional and hospital-type trends; a set of individual patient composition controls; an indicator for elective strike days; and an interaction between the elective strike day indicator and exposure. (5) The sample sizes differ from the descriptives in Table 2 because singleton groups of fixed effects are dropped from the estimation. This does not change the coefficients.

(as opposed to our more conservative assumption that composition did not change differentially across more or less exposed hospitals).[28] However, Table 8 shows that there was a substantial fall in the number of emergency arrivals and admissions on all-out strike days (but not differentially across more or less exposed hospitals), raising concerns that unobservable patient composition may have changed at the national level. As a result, we restrict our attention to comparing changes in outcomes across differently exposed hospitals.

6. Results

We first present our baseline results. We then present results for different patient groups to better understand who is affected by the strikes. Finally, we undertake a number of robustness checks.

6.1. Baseline results

Table 10 presents our estimated impact of strike exposure on emergency admissions. In the first column, the outcome is a dummy variable for whether the patient has an emergency readmission within 30 days of their original admission. In the second column, the outcome is a dummy variable for whether the patient dies in hospital within 30 days of their original admission.

The table shows that there was no statistically significant difference in either measure of patient outcomes between hospitals that were more or less exposed to the strikes. This suggests that hospitals were able to mitigate many of the negative consequences of the strikes.

As noted in Section 2.2, hospitals could react to the strike in several ways, including the use of temporary staff to treat patients, and postponing planned care in order to reallocate staff resources (such as consultants or nurses) towards treating emergency patients whose treatment could not be delayed. Unfortunately, the granular data required to study these actions in detail is unavailable: for example, there are no available data on day-to-day use of temporary staff, or the number of different staff groups present in the hospital on any given day.[29] We therefore cannot study the impact of most of these hospital decisions directly and our results should be interpreted as the effects of exposure to the strikes including all of the short-term responses that hospitals took to mitigate the effects of the strikes.[30]

An exception to this is that we can observe reductions in elective admissions on strike days. Table 8 shows that more exposed hospitals did reduce elective activity on strike days by more than less exposed hospitals, suggesting that this was

[28] In this case, β_2 could be interpreted as the impact of the strike if a hospital had no junior doctors, and the total effect of the strikes for each hospital j would be given by $\beta_2 + \beta_4 Exp_j$.

[29] HES does assign a unique consultant identifier to every episode of care, and this can be used to count the number of unique consultants assigned to patients on any given day. If we use the number of elective consultants recorded in HES each day as the dependent variable in Eq. (4), we see that a 10 percentage point increase in strike exposure was associated with a 5.3% fall in the number of consultants treating elective patients on strike days. These consultants could have been reassigned to treat emergency patients, but we do not find any statistically significant increase in the number of consultants recorded as treating emergency patients when we repeat the analysis for emergency consultants. However, this may be due to reassigned consultants fulfilling the roles of junior doctors and not being recorded as the consultant responsible for the patient in the HES records. Full results are available upon request.

[30] This also includes a possible scenario where staff at hospitals with different staff composition react differently to the strike (e.g. if junior doctors at hospitals where they account for a greater fraction of the doctor workforce are less likely to strike). No data are available on hospital-specific strike rates, with NHS England reporting that these data were not collected or held on a hospital-level basis in response to a Freedom of Information request).

Table 11
30-day readmission rates by ethnicity.

	(1)
All Strike * Exposure	0.0137
	(0.0179)
All Strike * Exposure * Mixed	0.200
	(0.134)
All Strike * Exposure * Asian	−0.0116
	(0.0607)
All Strike * Exposure * Black	0.179*
	(0.0985)
All Strike * Exposure * Other	−0.0245
	(0.142)
All Strike * Exposure * Unknown	−0.0328
	(0.0626)
Outcome Mean	0.152
N	31,786,339

Note:(1) Standard errors clustered at the hospital level are displayed in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. (2) The outcome variable is whether the patient has an emergency admission within 30 days of their original admission. (3) The regression includes hospital fixed effects; a set of time fixed effects; regional and hospital-type trends; a set of individual patient composition controls; an indicator for elective strike days; an interaction between the elective strike day indicator and exposure; a dummy variable for each ethnicity; the interaction between a dummy variable for each ethnicity and exposure; the interaction between a dummy variable for each ethnicity and the dummy variables for each type of strike days; and the interaction between a dummy variable for each ethnicity, dummy variables for each type of strike days and exposure. (4) The sample size differs from the descriptives in Table 2 because singleton groups of fixed effects are dropped from the estimation. This does not change the coefficients.

a strategic action that hospitals took to help mitigate the impact of the strikes. Cancelling elective procedures on strike days would both reduce the total number of patients requiring care in hospital, and potentially allow staff to be redeployed towards treating emergency patients. If cancellations did reduce the impact of the strike on emergency patients, then we would expect the inclusion of the estimates of the cancellations in our regressions to increase the magnitude of the estimated impact of strike exposure on emergency patient outcomes. However, Appendix Table D4 shows that although our estimated effect of strike exposure on readmissions is around 80% larger when controlling for estimated daily cancelled elective activity, it is still not statistically significant.[31] This suggests that equally exposed hospitals with different levels of elective cancellations did not deliver vastly different outcomes for emergency patients on strike days.

6.2. Who was most affected by the strike?

In the previous section we showed that, on average, patients in hospitals that were more exposed to the strikes experienced no worse outcomes on strike days. Below, we examine further whether there was any variation in the impact of being admitted to a more exposed hospital across patient groups on the basis of their demographic characteristics. Appendix Table A2 presents the numbers of patients in each of the groups analysed below.

We first examine differences in the impact of exposure to the strike across patients of different ethnicity. To do this, we split patients into six ethnicity groups: white, mixed, Asian, black, other and unknown.[32] We then re-estimate Eq. 4 including dummy variables for these groups, additional interaction terms for these groups with the strike dummies and strike exposure, and a three-way interaction between ethnicity, strike dummy and strike exposure.[33] We report the coefficients from this three-way interaction as the coefficient of interest: this can be interpreted as the additional effect of strike exposure on patient outcomes for patients of different ethnicities relative to white patients.

Table 11 shows the results. It shows that black patients experienced an increase in readmissions on all-out strike days in more exposed hospitals relative to white patients: an additional 10 percentage point increase in junior doctors relative to senior doctors in the pre-strike workforce led to an additional 1.8 percentage point increase in readmissions for black patients on strike days compared to white patients in the same hospital.[34] This difference is significant at the 10% level.

[31] The point estimate on all out strike days by exposure increases relative to our primary results (from 0.0136 to 0.0246). This is the direction we would expect if elective cancellations reduced pressure on hospitals and allowed them to focus more resources on emergency patients, but neither estimate is statistically significantly different from zero, or from one another.

[32] 7.3% of patients are reported with unknown ethnicity.

[33] This is a classic triple-difference regression with groups defined by ethnicity, treatment (strike day) and exposure to the strike.

[34] Appendix Tables C1 and C2 show there were no statistically significant changes in the composition of either of these groups by exposure.

Table 12
30-day readmission rates by deprivation.

	(1)
All Strike * Exposure	−0.0180
	(0.0345)
All Strike * Exposure * Q2	0.0232
	(0.0435)
All Strike * Exposure * Q3	0.0365
	(0.0416)
All Strike * Exposure * Q4	0.0599
	(0.0436)
Outcome Mean	0.152
N	31,425,765

Note: (1) Standard errors clustered at the hospital level are displayed in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. (2) The outcome variable is whether the patient has an emergency admission within 30 days of their original admission. (3) The regression includes hospital fixed effects; a set of time fixed effects; regional and hospital-type trends; a set of individual patient composition controls; an indicator for elective strike days; an interaction between the elective strike day indicator and exposure; a dummy variable for each deprivation group; the interaction between a dummy variable for each deprivation group and exposure; the interaction between a dummy variable for each deprivation group and the dummy variables for each type of strike days; and the interaction between a dummy variable for each deprivation group, dummy variables for each type of strike days and exposure. (4) The sample size differs from the descriptives in Table 2 because not all patients have a recorded address and singleton groups of fixed effects are dropped from the estimation.

The average readmissions rate for black patients over our whole sample period is 16.3%. A 10 percentage point increase in exposure therefore led to an 11.0% increase in readmissions for black patients relative to their average readmissions rate.

There were no other statistically significant differences in readmissions between other groups and white patients, although the coefficient for mixed ethnicity patients is of a slighty larger magnitude but less precisely estimated than for black patients. This lower precision could be driven by differences in sample size: the number of mixed ethnicity patients is 39% of the number of black patients (Appendix Table A2). The lack of statistically significant worse outcomes for mixed ethnicity patients admitted to more exposed hospitals should therefore be interpreted with care. Appendix Table D1 shows that when we combine the black and mixed ethnicity groups together, the point estimate is very similar to the coefficient for black patients only, and now significant at the 5% level.

The second dimension of patient heterogeneity we consider is local area deprivation. We assign patients to a deprivation quartile based on the Index of Multiple Deprivation (IMD) score of their local area in 2010.[35] The IMD score is the government's official measure of deprivation based on seven domains: income, employment, education, health, crime, barriers to housing and services, and living environment. Areas with a higher score (or in a higher quartile) are more deprived. We again adapt Eq. (4) to a triple-differences approach, with additional interaction terms between each deprivation group, strike exposure and strike drummies. We report the coefficients from this triple interaction, using the least deprived quartile as the base group.

Table 12 shows that there was no statistically significant difference in the effects of strike exposure for patients from different deprivation quartiles relative to the least deprived quartile. Taken together, these results show that while most patients were not affected by exposure to the strikes, there were meaningful impacts of exposure for some black patients, but no such impact of exposure for patients living in more deprived areas.

Our results therefore provide important evidence that the effects of staffing shocks are unlikely to affect all patient groups equally. The differences in outcomes across ethnicity groups are consistent with previous findings from the US, where despite having a very different healthcare system, minority groups have been shown to experience worse emergency care during periods of time when hospitals are particularly under strain (Singh and Venkataramani, 2022).

While we are unable to determine the exact mechanisms behind these differential effects, we can eliminate several potential explanations. First, our analysis includes hospitals fixed effects which absorb permanent differences between hospitals. This means that these differences in impacts of the strikes occur *within* hospitals, and so are not explained by differential exposure across patient groups to striking doctors. Second, our analysis also controls for a number of demographic and medical characteristics, which suggests that these differences are not driven by differences in clinical need. Finally, changes in readmission rates during strike days may reflect readmissions for postponed care, rather than unexpected admissions due to complications caused by lower care quality as a result of the strikes. To explore this further, Appendix Tables D2 and D3 repeat the analysis for length of stay and the number of recorded procedures; if treatment is being postponed then

[35] The local area is measured at the lower layer super output area (LSOA), a census unit with an average population of 1600.

Table 13
Patient outcomes with additional severity controls.

	(1) Readmissions	(2) Mortality
All Strike	−0.00126	0.000650
	(0.00798)	(0.00419)
All Strike * Exposure	0.0130	0.000518
	(0.0158)	(0.00849)
Outcome Mean	0.152	0.0373
N	31,634,448	31,634,448

Note: (1) Standard errors clustered at the hospital level are displayed in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. (2) The outcome variable in the first column is whether the patient has an emergency admission within 30 days of their original admission. (3) The outcome variable in the second column is whether the patient dies in hospital within 30 days of their original admission. (4) All regressions include hospital fixed effects; a set of time fixed effects; regional and hospital-type trends; a set of augmented individual patient composition controls; an indicator for elective strike days; and an interaction between the elective strike day indicator and exposure. (5) The sample sizes differ from the descriptives in Table 2 because singleton groups of fixed effects are dropped from the estimation. This does not change the coefficients.

we would expect to see a fall in length of stay and number of procedures on strike days, with patients returning for this treatment at a later data. Overall, there is no change in the average length of stay or number of treatments received by patients admitted to hospitals that were more exposed to the strike. There is also no differential change by exposure for black patients. These results suggest that the higher readmission rates for black patients at more exposed hospitals are not being driven by a reduction in the *quantity* of care provided to these patients as a result of the strike.

At least two potential explanations remain. First, black patients may be more likely to be treated by junior doctors, either due to unobserved differences in clinical need, or for other non-clinical reasons. If this was the case, then their treatment would be relatively more disrupted on strike days when fewer junior doctors are available. Second, the alternative arrangements for care made by hospitals on strike days may be less effective for treating black patients (e.g. if senior doctors are less effective than junior doctors at communicating with black patients). However, without data to link individual doctors to individual patients we are unable to distinguish between these two potential channels.

6.3. Robustness

In this section we test the robustness of our results to different specifications. One threat to identification is that, despite the evidence to the contrary in Section 4.2.2, patient composition is changing differentially on strike days across hospitals with different shares of junior doctors. Changes in patient outcomes may therefore simply reflect changes in patient characteristics that are not related to the quality of care: for example, if hospitals with a greater share of junior doctors treated relatively sicker patients on the strike days then we would overestimate the negative impacts of the strikes on patient outcomes.

To address this potential threat, we repeat our analysis with additional patient control variables. This includes interactions between age dummies, sex, and Charlson score fixed effects, and interactions between age, sex and primary diagnosis fixed effects. Table 13 shows the results for all patients and Table 14 shows the result split by ethnicity.

The results in Table 13 again confirm the null results found in the baseline specification when looking across all patients. This provides more evidence that our results are not being driven by changes in patient severity. The estimates in Table 14 further strengthen the evidence of differential effects of strike exposure between black and white patients: with the additional controls, the coefficient is slightly larger than that presented in Table 11 and is now statistically significant at the 5% level.

One remaining concern with our methodology, and in particular with the long sample period we use, is that more and less exposed hospitals may have differential trends in patient volumes or outcomes. Our estimates for the effects of exposure on strike days are relative to national and regional time trends, and so differential trends would bias our results. For example, if more exposed hospitals had a trend of worsening patient outcomes over the sample period, the estimated effect of exposure on strike days would be biased downwards. If this is the case, using a much shorter time period would substantially change our results.

To test this, Appendix Tables E1 and E2 show the results when we include the patient composition controls in our main specification but reduce the sample period to the eight months between November 2015 and June 2016 (the four months of the strikes and two months before and after). The results are qualitatively unchanged from our main results. This suggests that our results are not driven by our choice of sample period or differential time trends between more and less-exposed hospitals.

Table 14
30-day readmission rates by ethnicity with additional severity controls.

	(1)
All Strike * Exposure	0.0132
	(0.0188)
All Strike * Exposure * Mixed	0.203
	(0.133)
All Strike * Exposure * Asian	−0.0100
	(0.0618)
All Strike * Exposure * Black	0.218**
	(0.102)
All Strike * Exposure * Other	−0.00174
	(0.134)
All Strike * Exposure * Unknown	−0.0437
	(0.0638)
Outcome Mean	0.152
N	31,634,448

Note: (1) Standard errors clustered at the hospital level are displayed in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. (2) The outcome variable is whether the patient has an emergency admission within 30 days of their original admission. (3) The regression includes hospital fixed effects; a set of time fixed effects; regional and hospital-type trends; a set of individual patient composition controls; an indicator for elective strike days; an interaction between the elective strike day indicator and exposure; a dummy variable for each ethnicity; the interaction between a dummy variable for each ethnicity and exposure; the interaction between a dummy variable for each ethnicity and the dummy variables for each type of strike days; and the interaction between a dummy variable for each ethnicity, dummy variables for each type of strike days and exposure. (4) The sample size differs from the descriptives in Table 2 because singleton groups of fixed effects are dropped from the estimation. This does not change the coefficients.

A final test of whether our methodolgy is capturing the true effects of strike exposure or other trends in patient outcomes is to use placebo tests for non-strike days.[36] Appendix Tables E3–E10 show the results when we replace $AllStrike_t$ with a dummy variable for non-strike days. In particular, we use the days one and two weeks before and after the true dates of the all-out strike. In all but one case, the coefficient for black patients on exposure is insignificantly different from zero, and one is negative (i.e. the opposite sign of our primary result) and significant at the 5% level.

7. Conclusion

This paper examines the consequences of a nationwide junior doctor strike that took place in England in 2016. We exploit differences in the share of the doctor workforces who were eligible to strike across hospitals to estimate the impact of greater exposure to the strike on patient outcomes. We show that overall, higher exposure to the strike did not result in a higher rate of emergency readmissions or inpatient mortality over a 30 day period. While we were unable to determine the most effective actions that hospitals took to avoid harms to patients, our findings that more exposed hospitals reduced elective care the most demonstrates that hospitals could and did respond to the strike in a strategic way. This is encouraging, and suggests that hospitals were, in general, able to mitigate much of the impact of this strike on emergency patients.

More worrying, our results do however provide some evidence that there were differential effects of strike exposure for black and white patients. While our methodology means that we cannot estimate the total impact of the strikes on any patient group - due to national changes in patient volumes and composition on strike days - we show that black patients were more negatively affected by exposure to the strikes than white patients in the same hospitals. These differences are not found for patients of other ethnicities, or across more and less deprived groups. While we can rule out that these results were driven by differential changes in the (observed) characteristics of black and white patients treated in the same hospitals, and by changes in the quantity of care given to these groups, the ultimate mechanisms underlying these differences remain unclear. Better understanding of these mechanisms - and ultimately how to prevent such outcomes recurring in future - should be an important priority for future research.

Our results provide important evidence on the effects of physician strikes. Unlike much of the existing literature on medical strikes, we are able to ensure that our results are not driven by compositional changes and can isolate the patient groups that were most affected. Real-terms declines in average pay over the last decade and COVID-related pressures mean

[36] This is also similar in nature to an event-study approach testing for pre-trends as it shows how outcomes changed on the same day of the week in the weeks preceding (and following) the strike

that British medical unions are once again taking industrial action[37], and this paper provides important evidence as to the potential impact of such action.

Our results also provide important information about the impacts for different patient groups when hospitals operate at or near capacity. Previous work has shown that unexpected demand shocks which put hospitals under strain can lead to worse outcomes (Evans and Kim, 2006; Hoe, 2022; Singh and Venkataramani, 2022). We examine what happened to patient outcomes during a temporary supply shock. The absence of significant negative effects on aggregate suggest that NHS hospitals could temporarily react to mitigate many of the worse outcomes, although some groups still experienced worse outcomes.

One limitation of our work has been that we have not been able to study in detail which decisions or actions undertaken by hospitals were most successful in preventing harm to patients. This is important for planning responses to future strikes in the UK and in other countries, and in deciding whether such actions are sustainable in the long-term. Some of these actions would not be available to hospitals in other countries, and may now be harder to take in the UK than in 2016. For example, postponing non-urgent activity is a clear action that NHS hospitals took to reallocate resources during the strike. However, this still causes disruption to patients, and may be more difficult to achieve going forward when waiting lists are already much longer following the COVID pandemic. Temporarily drafting in replacement staff - either by reallocating other hospital staff to cover tasks usually performed by junior doctors, or by using staff from primary care and ambulance providers - is also likely to be an important response to any strikes, but may be harder in NHS hospitals if these staff are less able or less willing to fulfil these roles than they were in the past. This may also be very difficult to achieve in other countries where healthcare systems are less integrated. Given the substantial pressures on the hospital workforce around the world, both in the short-term as a result of the COVID pandemic, and in the long-term given rising demands on healthcare workers, further research into such mechanisms remains a priority.

Declaration of Competing Interest

None.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at 10.1016/j.jebo.2023.06.011.

References

Chima, S.C., 2020. Global medicine: is it ethical or morally justifiable for doctors and other healthcare workers to go on strike? BMC Med. Ethics 14 (S5).
Costa, E., 2022. The unintended consequences of hospital strikes on patient outcomes evidence from multiple strikes in the Portuguese national health service. Health Econ. 31 (11), 2499–2511.
Cunningham, S.A., Mitchell, K., Narayan, K.V., Yusuf, S., 2008. Doctors' strikes and mortality: a review. Social Sci. Med. 67 (11), 1784–1788.
Erceg, M., Kujundžić-Tiljak, M., Babić-Erceg, A., Ćoric, T., Lang, S., 2007. Physicians' strike and general mortality: Croatia's experience of 2003. Collegium Antropologicum 31 (3), 891–895.
European Foundation for the Improvement of Living and Working Conditions, 2011. Employment and industrial relations in the health care sector. Last accessed 28 January 2022. https://www.eurofound.europa.eu/sites/default/files/ef_files/docs/eiro/tn1008022s/tn1008022s.pdf.
Evans, W.N., Kim, B., 2006. Patient outcomes when hospitals experience a surge in admissions. J. Health Econ. 25 (2), 365–388.
Friedman, W., Keats, A., Kavao Mutua, M., 2022. Disruptions to healthcare quality and early child health outcomes: evidence from health-worker strikes in Kenya. Am. Econ. J. Econ. Policy 14 (2), 231–262.
Furnivall, D., Bottle, A., Aylin, P., 2018. Retrospective analysis of the national impact of industrial action by English junior doctors in 2016. BMJ open 8 (1), e019319.
Gruber, J., Kleiner, S.A., 2012. Do strikes kill? Evidence from New York state. Am. Econ. J. Econ. Policy 4 (1), 127–157.
Health Education England, 2017. Junior doctor morale: understanding best practice working environments. Last accessed 3 November 2021.
Hoe, T.P., 2022. Does hospital crowding matter? Evidence from trauma and orthopedics in England. Am. Econ. J. Econ. Policy 14 (2), 231–262.
James, J.J., 1979. Impacts of the medical malpractice slowdown in los angeles county: January 1976. Am. J. Public Health 69 (5), 437–443.
NHS England, 2016a. NHS England update on industrial action by junior doctors. Last accessed 8 October 2021. https://www.england.nhs.uk/2016/01/junior-doctors/.
NHS England, 2016b. NHS England update on industrial action by junior doctors. Last accessed 8 October 2021. https://www.england.nhs.uk/2016/02/junior-doctors-update/.
NHS England, 2016c. NHS England update on industrial action by junior doctors. Last accessed 8 October 2021. https://www.england.nhs.uk/2016/03/update-junior-doctors/.
NHS England, 2016d. NHS England update on industrial action by junior doctors. Last accessed 8 October 2021. https://www.england.nhs.uk/2016/04/update-junior-doctors-2/.
NHS England, 2016e. NHS England update on industrial action by junior doctors. Last accessed 8 October 2021. https://www.england.nhs.uk/2016/04/update-junior-doctors-3/.
NHS England, 2016f. NHS England update on industrial action by junior doctors. Last accessed 8 October 2021. https://www.england.nhs.uk/2016/04/update-junior-doctors-4/.
Quan, H., Sundararajan, V., Halfon, P., Fong, A., Burnand, B., Luthi, J.-C., Saunders, L.D., Beck, C.A., Feasby, T.E., Ghali, W.A., 2005. Coding algorithms for defining comorbidities in ICD-9-CM and ICD-10 administrative data. Med. Care 1130–1139.
Roemer, M.I., Schwartz, J.L., 1979. Doctor slowdown: effects on the population of los angeles county. Social Sci. Med. Part C Med. Econ. 13 (4), 213–218 .

[37] https://www.theguardian.com/society/2023/apr/11/junior-doctors-in-england-begin-four-day-strike-over-pay

Ruiz, M., Bottle, A., Aylin, P., 2013. A retrospective study of the impact of the doctors' strike in England on 21 June 2012. J. R. Soc. Med. 106 (9), 362–369.
Salazar, A., Corbella, X., Onaga, H., Ramon, R., Pallares, R., Escarrabill, J., 2001. Impact of a resident strike on emergency department quality indicators at an urban teaching hospital. Acad. Emerg. Med. 8 (8), 804–808.
Siegel-Itzkovich, J., 2000. Doctors' strike in israel may be good for health. BMJ 320 (7249), 1561. doi:10.1136/bmj.320.7249.1561.
Singh, M., Venkataramani, A., 2022. Capacity Strain and Racial Disparities in Hospital Mortality. NBER Working Paper 30380.
Slater, P., Ever-Hadani, P., 1983. Mortality in jerusalem during the 1983 doctors'strike. The Lancet 322 (8362), 1306.