

NBER WORKING PAPER SERIES

CAPACITY CONSTRAINTS AND THE PROVISION OF PUBLIC SERVICES:
THE CASE OF WORKERS IN PUBLIC HEALTH CLINICS

Matthew C. Harris
Yinan Liu
Ian McCarthy

Working Paper 25706
<http://www.nber.org/papers/w25706>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2019

We are grateful to the Knox County Health Department, particularly Dr. Kelly Cooper, Natalie Alverson, and Caley Webster for their partnership and use of their data. We are also grateful for comments from Bill Neilson, Sara Markowitz, Scott Barkowski, David Bradford, Don Bruce, Georg Schaur, Don Clark, and participants at iHEA 2017 and the 2017 Southeastern Health Economics Association Meetings. All errors are our own. 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Matthew C. Harris, Yinan Liu, and Ian McCarthy. 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.

Capacity Constraints and the Provision of Public Services: The Case of Workers in Public Health Clinics

Matthew C. Harris, Yinan Liu, and Ian McCarthy

NBER Working Paper No. 25706

March 2019

JEL No. D24,I18,J22

ABSTRACT

Unlike in the production of most goods, changes in capacity for labor-intensive services only affect outcomes of interest insofar as service providers change the way they allocate their time in response to those capacity changes. In this paper, we examine how public sector service providers respond to unexpected capacity constraints in the specific context of public health clinics. We exploit an exogenous reduction in public health clinic capacity to quantify nurses' trade-off between patients treated and time spent with each patient, which we treat as a proxy for a quality v. quantity decision. We provide evidence that these small and generally insignificant effects on nurse time favor public sector employees prioritizing quality of each interaction over clearing the patient queue.

Matthew C. Harris
University of Tennessee
Haslam College of Business
722A Stokely Management Center
916 Volunteer Boulevard
Knoxville, TN 37996
mharris@utk.edu

Ian McCarthy
Department of Economics
Emory University
Rich Memorial Building, Room 319
Atlanta, GA 30322
and NBER
immccar@emory.edu

Yinan Liu
1640 Cumberland Ave.
Howard H. Baker Jr. Center for Public Policy
University of Tennessee
Knoxville, TN 37996
yliu110@vols.utk.edu

1 Introduction

The public labor force is a large component of the economies of most developed countries. Over 15 percent of the labor force is employed in the public sector in the U.S., and payroll expenses account for approximately half of all state government expenditures (OECD, 2017; Gordon and Iselin, 2017). Many of these workers are engaged in the production of labor-intensive services. For example, according to the Bureau of Labor Statistics, six of the ten most common occupations for government workers include: clerks, postal service workers, repair workers, highway workers, corrections officers, and (most importantly for this paper) registered nurses.

Most service providers, in both the public and private sector, face inherently stochastic demand but cannot store inventory. Suppliers therefore tend to carry excess capacity on a median day (DeVany, 1976); however, the presence of excess capacity is heavily influenced not only by stochastic fluctuations in demand but also variations funding. Such variation in capacity constraints due to funding is a particularly salient issue for public services, where debates over funding levels and threats of budget cuts are regular occurrences with important consequences. For example, DeAngelo and Hansen (2014) show that when budget cuts forced layoffs of State Troopers in Oregon, traffic fatalities increased. Other work examining law enforcement show that capacity constraints do affect the provision of public services, providing evidence that greater funding for law enforcement decreases crime (Evans and Owens, 2007; Chalfin and McCrary, 2018; Mello, 2019).

Given that capacity constraints affect outcomes related the provision of public services, how do employees respond when capacity constraints bind and prices cannot adjust?¹ Understanding how providers of public services reallocate their time when demand exceeds capacity (or vice versa) is critical for understanding how changes in funding for the provision of public services will affect outcomes of interest. In a for-profit setting, we expect that employees will respond in some way consistent with the profit maximization of the firm, or respond to the incentives in place to alleviate principal agent problems. However, there is a large body of work showing that employees in non-

¹Prior work has examined the effects of congestion (i.e., capacity constraints) under different levels of demand in the transportation industry, and when prices are efficient allocation mechanisms (Carlin and Park, 1970; Daniel, 1995; Brueckner, 2002; Mayer and Sinai, 2003).

profit and public settings fundamentally differ from workers who select into for-profit employment, making the response of public sector employees *ex ante* less clear (Preston, 1989; Leete, 2001; Benz, 2005; Mas, 2006; Prendergast, 2007; Delfgaauw and Dur, 2008; Buurman et al., 2012).

In this paper, we provide evidence on how service workers in the public sector respond to capacity constraints, focusing specifically in the context of public health clinics. We examine how nurses in public health clinics respond to unexpectedly tight time constraints created by exogenous temporary reductions in staff. When demand spikes and capacity is constrained, how much time with the average patient are providers willing to trade-off to see as many patients as possible? While the specific answer within this context can inform about the amount of median excess capacity built into the provision of public clinical services, more broadly this paper examines how providers of public services make quality v. quantity tradeoffs under binding time constraints. For example, if providers are reluctant to reduce the time spent with patients, and leaving patients untreated has a high social cost, then ensuring median day excess capacity in the public provision of clinical services can yield significant welfare gains (Hughes and McGuire, 2003). While leaving patients untreated creates obvious negative externalities, particularly if they are in the clinic for communicable illness, a reduction in time spent with each patient may also negatively affect the quality of care.²

We identify the causal effects of reductions in capacity through a series of repeated, but not periodic, exogenous temporary reductions in the number of nurses working in a given clinic on a given day. Our data were provided by the Knox County Health Department (KCHD) in Tennessee and are comprised of time records for each patient visit in five public health clinics over sixteen months. In addition to providing certain types of health care in the clinic, KCHD is also responsible for administering FluMist vaccines in public schools. On days when KCHD is administering FluMist, two nurses would be removed from typical clinical duties and sent to the particular school for the morning, leaving clinics short-staffed with reduced operating capacity for the first half of the day.

²See, for example, Whittington and McLaughlin (2000), Tai-Seale et al. (2007), Munyisia et al. (2011), and McCloskey et al. (2014), among others.

The selection and timing of FluMist administration is plausibly exogenous to the demand or expected patient volume for a given clinic. For example, all scheduling decisions were made by the KCHD central office without consulting the clinics and with no compensating actions taken by the central office. Clinics that were selected for FluMist on a given day were instructed to keep all scheduled visits and were prohibited from otherwise increasing their staffing levels on FluMist days. Indeed, we were asked to examine these data by KCHD because the effects of FluMist administration on clinical production were unknown. KCHD wanted to know if their current practices in conducting FluMist had any adverse effects on their clinical mission. We therefore contend that FluMist-induced staffing shortages are exogenous to the scheduled daily activity of a given clinic, and the number of their scheduled patients was not influenced by the staff shortage. We expand on these institutional details and provide empirical evidence on the exogeneity of FluMist days in [Section 2](#).

Our empirical analysis exploits these reductions in clinic staff, along with unexpectedly high-demand days, to quantify a provider’s trade-off between patients seen versus time with each patient. Our analysis proceeds as follows. First, we quantify the effects of reductions in clinic staff on provider behaviors and specific components of the clinic visit. A simple event study of daily clinic visits and other aggregate measures of clinic behaviors surrounding FluMist days shows an abrupt reduction in capacity on FluMist days. Our regression analysis of provider-days further confirms that nurses removed for FluMist administration see significantly fewer patients, thereby reducing overall clinic capacity. We also find that providers in affected clinics (who are not administering FluMist) decrease their share of walk-in patients, indicating they are prioritizing those patients with scheduled appointments. At the visit level, when clinic capacity is reduced, average total visit time significantly decreases by 7% (or about 5 minutes). This primarily occurred through a reduction in check-in and check-out times, with small (and insignificant) reductions in time with nurses.

Next, we consider the underlying mechanisms that may drive our estimated effects. To guide our analysis, we construct an expository theoretical model of providers’ responses to staff reductions, where we posit that a provider’s utility is a function of the number of patients seen and the

amount of time spent with each patient, relative to some threshold ‘sufficient’ visit length. We derive comparative statics showing that the optimal amount of time that providers spend with patients is a function of the relative importance of visit length versus number of patients seen and the stochastic arrival rate of patients, among other parameters. In the context of our theoretical model, the null effects on provider time with patients are reflective of providers’ preference for time spent with each patient over the number of patients seen. We then conduct additional analyses to test whether our results are plausibly driven by provider preferences or simply a reflection of existing excess clinic capacity. For example, we estimate unconditional quantile regressions allowing for differential effects along the support of daily visit volumes. Even on the busiest days (upwards of the 75th percentile of visit volume) when capacity constraints are likely binding, nurses never reduce their time with each patient by more than 5%. We interpret the inelasticity of time spent with patients as indicative of provider preferences; however, we acknowledge that our findings may be partially driven by the structural constraints of the provider-patient interaction.³

Focusing on a specific institution and context allows for a strong identification strategy and thus aids our causal analysis. This ultimately improves the internal validity of our analysis, but potentially at the expense of generalizability. Nonetheless, we contend that our results are at least partially generalizable to other public provided services, particularly those where demand is stochastic. For example, during the government shut down in the winter of 2018-2019, Transportation Security Administration (TSA) employees called out sick, and public reports suggested significantly longer queues for airline passengers. TSA officials therefore appear to have made the choice to maintain (or at least approximate) pre-existing screening standards rather than more quickly clear the queue of passengers. Additional areas where employees engaged in the provision of services to the public may exhibit similar responses include postal service, guidance counselors in public schools, public defenders, and rehabilitation facilities. While some of these entities are more leanly funded than others, the effects of budget cuts or increases to any of these services depends on how providers manage tradeoffs between quality of service and customers served. Our

³In other words, providers may truly be unable, rather than unwilling, to shorten visit lengths to clear the waiting room. In either case, the empirical and policy implications of our results for reducing clinic capacity are the same.

analysis and conceptual framework may offer insights into future studies in these other important areas.

Our study offers three distinct contributions to the literature. First, our investigation of how public service workers adjust their time allocation in response to reductions in capacity is novel.⁴ Most prior relevant work in the health sector focuses on excess capacity and provider response to stochastic demand in the hospital and long-term care settings (Friedman and Pauly, 1981; Gaynor and Anderson, 1995; Keeler and Ying, 1996; Hughes and McGuire, 2003; Sharma et al., 2008) or providers' acceptance of patients and time spent with each patient in emergency departments (Chan, 2016, 2018). Recent work from Freedman et al. (2018) examines changes in physician behavior due to increased time pressures in the clinic, exploiting variation in patient volumes to identify responses of primary care physicians within the day. Our analysis is similar in spirit, albeit with a different source of identification (a reduction in nurse staffing levels in the clinic) and a different care setting (public vs. private). There are also reasons to suspect that workers in the public and private sectors may respond very differently to reduced capacity. For example, Dixit (2002) discusses how incentives and competition can inefficiently distort worker effort and performance in the public sector.⁵

Second, most prior work on exogenous capacity changes in health care settings focuses on increased capacity rather than reduced capacity. For example, there have been a number of studies that exploit regulation changes in required staffing/patient ratios as exogenous shocks to staffing levels and investigate the effects of the regulation change, with mixed findings.⁶ In addition, previous studies that prompted such regulation change have been criticized for problems

⁴In studies of other industries, understaffing has been found to be related to lower levels of performance at the group level in professional and trade occupations (Ganster and Dwyer, 1995), a decline in the positive experiences and increased workload stress in an educational service setting (Yoe, 1988), and less than optimal sales and profitability in stores (Mani et al., 2015).

⁵The public health setting is important in its own right as over 20 million people currently receive primary and preventative health care at community health centers (Kaiser Family Foundation, 2013). Additionally, capacity constraints may have differential effects when the constraint is on labor, rather than capital (beds), or when the need for treatment is more/less urgent. Unlike emergency departments, most patients to public health clinics will survive until the next day if untreated, in which case providers in health clinics may place more weight on time with patients over maximizing the number of patients seen in a timely manner.

⁶Chen and Grabowski (2015), Bownblis (2011), Park and Stearns (2009), Tong (2011), Aiken et al. (2010), and Lin (2014) found quality of care increased in at least one dimension, while Evans and Kim (2006), Matsudaira (2014), and Cook et al. (2012) found no change in quality of care.

including omitted variable bias and endogeneity of staffing levels (Evans and Kim, 2006).⁷ To our knowledge, this is the first study to investigate the effects of exogenously *decreased* staffing levels on time spent with patients and number of patients seen.

Finally, whereas prior work often examines *permanent* regulation-induced changes in staffing levels, we study the effects of temporary staffing decreases. For example, previous studies have linked ‘lower than target’ nurse staffing levels and higher patient turnover with higher mortality rate on a daily basis (Needleman et al., 2011; Schilling et al., 2010). Our results indicate that effects of staff reductions were strongest on days with the largest patient volume, which suggests that estimates derived from a permanent capacity change may mask larger effects on critical days.

2 Data and Institutional Details

Data were provided by the Knox County Health Department (KCHD) in Tennessee and are comprised of time records for each patient visit in five public health clinics over 16 months and two flu seasons. Each individual record was documented by clinic staff in an electronic patient record, where we observe the date of the visit, the initiation of the visit (scheduled or walk-in), the location (clinic) of the visit, the age range of the patient, and the unique provider/nurse ID for each visit. We also observed detailed time stamps for different stages of each visit, including: 1) Check-in time, the time between signing in and being taken to a treatment room; 2) Ready Nurse time, the time spent in the treatment room awaiting a nurse; 3) Nurse time, the time spent from the start of the consultation to the conclusion of any treatment; and 4) Ready Check-Out time, the time between the conclusion of treatment and when the patient leaves.

KCHD provides many services to the community, including health education, awareness, vaccinations, and clinical services. Clinical services in the KCHD health clinics, the focus of this paper, are provided almost exclusively by registered nurses (RNs) rather than physicians. In addition, KCHD administers FluMist vaccines to public school children in Knox County, typically in October, November, or December. On FluMist days, two RNs are pulled from a subset of the five

⁷Variation across hospitals that could not all be captured might contribute to quality of care, or patients admitted during the weekend tend to have more severe conditions than those admitted during the week.

main clinics to administer FluMist in schools, subsequently reducing capacity in the nurse’s clinic during that time. On a FluMist day, nurses on FluMist duty were away the whole morning and would return to work in the clinics in the afternoon.

In total, our data consist of 42,514 visits to five public health clinics from September 2014 through January 2016. Approximately 6% of our observed visits occurred on a FluMist day. Overall summary statistics are provided in Table 1, where we present statistics for all clinic visits in the first column and statistics by FluMist/Non-FluMist days in columns 2 and 3, respectively. Statistics at the clinic level are summarized in the top panel of Table 1, with statistics on individual components of each visit in panel 2 and general patient/visit characteristics in panel 3.

TABLE 1 HERE

As shown in panel 1, clinics saw around 25 patients on average per day, where approximately 33% of visits were scheduled and 67% were walk-ins. From panel 2 of Table 1, clinic visits last around 73 minutes on average, with shorter visit lengths of 66 minutes on FluMist days. Time spent with nurses is the most time consuming aspect of a visit, with average nurse times of around 30 minutes. Nurse time and ready check-out times were comparable on FluMist days relative to non-FluMist days, while check-in times and ready nurse times were shorter. Finally, panel 3 of Table 1 presents the percentage of patients in different age groups as well as the percentage of different reasons for the visit, the day of the visit, and the clinic. These statistics suggest that the age distribution across patients is similar between FluMist and non-FluMist days, as is the underlying reason for the patient visit. We also see that FluMist days are not isolated to specific days of the week or disproportionately concentrated among a single clinic.

2.1 FluMist Administration

Two key features of the administration of FluMist vaccinations are useful in establishing FluMist as an exogenous source of temporary reductions in clinic capacity. First, nurses pulled from the clinic to administer FluMist in schools were not replaced by nurses from other clinics or temporary staff. Second, all scheduling decisions of FluMist days were made by the KCHD central

office without consulting the clinic. When a clinic was selected for a FluMist day, the staff who remained were instructed to maintain their scheduling patterns and staffing levels. In other words, clinics that had RNs out at schools were told to treat the day like a normal day – but with fewer clinicians.⁸

While the institutional details of FluMist administration suggest that staffing reductions were indeed exogenous, it is of course possible that clinics anticipated the FluMist days and adjusted accordingly. We therefore consider a series of event studies that examine daily clinic patterns for the days immediately surrounding a FluMist day. By design, our event studies are not regression-based and instead reflect basic descriptive statistics over time. The purpose of these event studies is twofold: 1) to illustrate the reduction in capacity from FluMist administration; and 2) to examine whether clinics anticipated the staffing reductions in some way and adjusted their behaviors leading up to FluMist days. Results are summarized in Figure 1, where we present statistics for clinic visits (total, scheduled, and walk-in) and total minutes spent in each stage of a visit, including nurse-patient time, check-in time and check-out time for each day within two business days of a FluMist day.

FIGURE 1 HERE

Figure 1a depicts total daily visit volume at the clinic level for ± 2 business days surrounding a FluMist Day. For each of the two days before and after a FluMist day, the clinic sees an average of 25.8 patients. But on FluMist days, the clinic sees an average of 21.8 patients, which represents more than a 15% decrease in total patient volume. Figures 1b and 1c plot total daily visit volume for scheduled visits and walk-in visits, respectively. There is some descriptive evidence that providers in FluMist clinics may anticipate these FluMist days by scheduling (on average) one fewer appointment on FluMist days than neighboring days; however, clinics selected for FluMist see 3.5 fewer walk-in patients on a FluMist day than neighboring days. The fact that the reduction in

⁸It is also worth noting that there were no compensating actions taken in any way by the central office. We were asked to examine the visit level data from KCHD because the consequences of these short-staffing days with respect to quality of care or production of public health were not understood. The central office wanted to know what (if any) compensating actions should be taken.

total visits is primarily driven by decreased walk-in volume indicates that FluMist days do indeed reduce clinic capacity.⁹

We also examine how FluMist days differ from adjacent days in terms of total time patients spend with their providers and in administrative components of the visit such as check-in and check-out times. Figure 1d shows that on FluMist days, patients collectively spend a total of 570 minutes with nurses throughout the day compared to 725 minutes with nurses on adjacent days. Consistent with the overall reduction in total visits, this abrupt decrease in total nurse-patient minutes on FluMist days is reflective of a reduction in clinic capacity on those days.

In summary, these figures provide descriptive evidence that FluMist days reduce total clinic capacity. To the extent that providers are able to anticipate the effects of FluDays, they are only able to slightly adjust the demands placed on the clinic via small reductions in scheduled visits. Walk-in visits account for about two-thirds of total visit volume, and three-fourths of the drop in patients seen is due to decreases in walk-in patients. Nonetheless, these preliminary takeaways are purely descriptive and based on simple clinic means. In subsequent sections, we turn to regression methods to examine how nurses respond to reduced capacity and tighter time constraints when controlling for a rich set of patient, clinic, visit, and provider characteristics.

3 Initial Evidence on the Effects of Capacity Reductions

We first provide initial evidence of the average effect of FluMist days on various activities at both the provider and visit level, respectively. This analysis provides another reasonableness check that FluMist administration does indeed reduce clinic activity among those nurses directly impacted. We then further show the effects of FluMist days on nurses that remain in the clinic (i.e., nurses that were exposed to a reduction in capacity but were not removed from the clinic for FluMist administration). Details of these analyses and findings are discussed throughout the

⁹This comparison differs somewhat from the overall summary statistics in Table 1, where we see more scheduled visits on an average FluMist day and a comparable number of walk-in visits relative to an average non-FluMist day; however, the comparison group of days in this event study is only for days immediately surrounding FluMist days rather than all non-FluMist days over the sample period. The numbers reflected in Figure 1 therefore offer a more apples-to-apples comparison between FluMist and non-FluMist days.

remainder of this section.

3.1 Provider-level Analysis

For our provider-level analysis, we construct a panel of provider/days and estimate the following fixed effects model:

$$y_{it} = \alpha + \beta_n FluNurse + \beta_d FluDay + \mu_i + \nu_c + \eta_d + \gamma_m + \delta_y + \varepsilon_{it}. \quad (1)$$

We denote daily output for a given clinician (nurse) i at time t by y_{it} , measured as log numbers of scheduled visits, log number and share of walk-ins, and log time spent with patients. The variable *FluNurse* is an indicator set to one for a particular nurse if he or she was on FluMist duty on that day. Similarly, we form a *FluDay* indicator that takes a value of one if any nurse from that clinic administered FluMist on that day. Therefore, if provider i from clinic c is on FluMist duty on a given day, both the *FluNurse* and *FluDay* indicators are set to 1. Meanwhile, if some other provider ($\neg i \in c$) from i 's clinic is on FluMist duty, then *FluDay* will equal 1 but *FluNurse* will be 0. From the provider's perspective, the indicator for *FluDay* therefore implies an increase in the expected number of patients to be seen by each each provider who remains in the clinic on a FluMist day. We estimate this model using a fixed effects "within-estimator" at the nurse level, also including fixed effects for each clinic (ν_c), day of week (η_d), month of year (γ_m), and year (δ_y).¹⁰ Standard errors are clustered at the nurse level.

TABLE 2 HERE

Table 2 presents our provider-level estimates of the average effects of FluMist-induced staff reductions on nurses' daily production. Column (1) presents the estimated effect from being called out of the clinic to administer FluMist on a given day. These estimates are based on the full sample. Column (2) presents estimates on the effect of a FluMist day among nurses who were not removed from the clinic on that day. The estimates in column (1) of Table 2 therefore provide a

¹⁰Among other things, the inclusion of nurse fixed effects captures any potential selection at the clinic level with regard to which nurses are ultimately pulled from the clinic to administer the FluMist vaccine.

reasonableness check for our provider-level analysis, as these estimates reflect changes to behaviors specifically for nurses who are removed from the clinic to administer the FluMist vaccines. Since FluMist nurses typically spend a little less than half of their day out of the clinic, our estimates that total time spent with patients and total patients seen decreases by around 45% for FluMist nurses are in-line with *a priori* expectations. We also find that being scheduled for FluMist increases the nurse’s share of walk-in patients relative to scheduled patients. This is consistent with a backlog of walk-in patients on FluMist days, where nurses staying in the clinic prioritize scheduled patients over walk-in patients, and upon their return to the clinic, nurses out for FluMist administration work to alleviate the queue of walk-in patients.

The estimates in column (2) of Table 2 reflect the estimated effect of FluMist days on non-FluMist nurses (i.e., the nurses that remained in the clinic for the entire day). Here, we find a statistically significant but small increase in total visits (at the 90% confidence level) and a significant decrease in the share of walk-in visits (at the 95% level) among non-FluMist nurses, with the latter result again suggesting a prioritization of scheduled visits over walk-in visits.

These results are consistent with providers placing more value on spending a certain amount of time with each patient relative to seeing as many patients as possible. For example, the average clinic has around 6 nurses staffed in a given day. Typically, two nurses are removed for the morning to administer FluMist, leaving four nurses remaining in the clinic. The estimates for FluDay in column (2) of Table 2 suggest that the remaining nurses collectively increase patient time that day by about 0.25 average-person days (6.3 percent increase, presumably of a normal day’s activity $\times 4 \approx 0.25$ additional person days), see an additional 24 percent of a person-day’s equivalent of patients, and see an additional 38 percent of a person-day equivalent scheduled patients. Given that on a FluMist day, the clinic loses nearly a full person-day of capacity (and activity), the magnitudes of these increases do not compensate for the reduction from RNs removed from the clinic. Note also that while the nurses who remain in the clinic may see slightly more patients, they do not appear to be sacrificing average time with patients to do so. In summary, the sign of the coefficients on the FluDay indicator are consistent with some form of compensating behavior, but the estimates are often statistically insignificant and the magnitudes are insufficient to fully

compensate for the reductions in output from nurses temporarily removed from the clinic.¹¹

3.2 Visit-level Analysis

We also examine the effects of capacity reductions on the average time spent in each stage of a visit. We adopt a similar specification as in equation 1, with three main differences: 1) we include a larger set of fixed effects, including patient age (in 10-year bands), clinic, provider, reason for visit, day of the week, month, and year; 2) we only consider the FluDay indicator, since this indicator overlaps with the FluNurse indicator at the visit level; and 3) our visit-level outcome measures include total visit time, check-in minutes, waiting room time, nurse minutes, and check-out minutes (all in logs), as well as an indicator for whether the visit is a walk-in. Since patients do not visit clinics with sufficient frequency over time, we estimate our visit-level model using ordinary least squares. We specify the visit-level model as:

$$y_{vt} = \alpha + \beta FluDay + \mu_i + \nu_c + \eta_d + \gamma_m + \delta_y + \rho_v + a_v + \varepsilon_{vt}, \quad (2)$$

where arguments are defined as in equation 1 but with fixed effects for the reason for the visit (ρ_v) and age range of the patient (a_v).

Table 3 presents the estimated effects of FluMist on total visit time, time spent in different components of the visit, and the probability a visit is a walk-in. These results again indicate that providers value spending time with each patient over clearing all patients from the waiting room. Specifically, while we find a reduction in time spent in the waiting room, these estimates are imprecisely estimated. We also find a larger 8-10% reduction in the length of time spent in the check out process, and we estimate a slight reduction of 3% (significant only at the 90% level) in the length of time with a nurse; however, the effect on time with nurses appears to be driven by the nurses who are temporarily removed from the clinic for FluMist administration. Also, note that on FluMist days, visits are more than 10% less likely to be walk-in patients, implying

¹¹If capacity constraints are at all binding (even for just a portion of the time), some compensating behavior is to be expected. The job of these providers is to manage demands on the clinic as a whole rather than demand for their specific personal services.

that scheduled patients get priority when time constraints bind. Overall, patients’ total visit time on a FluMist day decreased by at least 7%, regardless of whether they were seen by a nurse who administered FluMist on that day, but this reduction is driven by streamlining administrative areas of the process, in particular check-out times, with no significant reduction in nurse minutes among non-FluMist nurses. Given that a FluDay represents, on average, a 16% reduction in production capability, the compensations we see are far from complete.

TABLE 3 HERE

4 Theoretical Framework

While the initial analysis in Section 3 shows some evidence of nurse behaviors when short-staffed, several features of the clinical context and the FluMist vaccine administration enable greater insight on the underlying mechanisms driving these results. In particular, what (if anything) do these results say about a provider’s preferences and willingness to trade off time with patients versus patients seen? To that end, we borrow elements from DeVany (1976) and Anand et al. (2011), among others, to motivate further empirical analysis with a hypothetical loss function for a provider engaged in the production of public services.

This conceptual framework is a natural fit for our research question as it accommodates two key stylized facts of services where labor is the primary input and quality of the service provided is a function of time with the customer. First, since demand is stochastic, clinics have some excess capacity on a median day. Second, given an expected arrival rate of patients to the clinic as a whole, a reduction in the number of providers is equivalent to a proportional increase in the arrival rate of patients to a remaining provider. The goal of the agents in our context is not to maximize profits but instead to ensure that a patient’s needs for care are met.¹²

Individuals (patients) are assumed to arrive following a Poisson process, with mean arrival rate, denoted λ , over a unit of time normalized to one. Service time is also assumed to be distributed

¹²Prices are pre-determined by KCHD, and most patients visiting the public health clinics face a nominal price of zero.

exponentially, with a mean service time denoted by μ . Providers minimize a loss function in each period (day) with respect to the average time spent with each patient (μ). We assume the function is additively separable in two arguments: 1) disutility from spending less time on average with patients than some fixed ideal amount of time, denoted by τ ; and 2) disutility from leaving patients unseen. Assuming that the mean service time is less than one, the number of unseen patients can be expressed as $(\lambda - 1/\mu)$, and the provider's loss function can be written as

$$U(\mu|\tau, \lambda) = f(\tau - \mu) + g(\lambda - \frac{1}{\mu}). \quad (3)$$

By definition, both $f(\cdot)$ and $g(\cdot)$ are assumed to be decreasing and convex. We impose a convenient functional form to derive a comparative static and evaluate how shocks to the provider's arrival rate, λ , affect the provider's optimal choice of time spent with patients. Assuming that $f(\cdot)$ and $g(\cdot)$ are exponential functions,

$$U(\mu|\tau, \lambda) = -e^{\alpha(\tau-\mu)} - e^{\beta(\lambda-\frac{1}{\mu})}, \quad (4)$$

where α captures the disutility from spending less time than ideal with patients, and β captures the disutility from leaving patients unseen. Taking the derivative with respect to μ yields the first order condition for the optimal amount of time spent with a patient, μ^* :

$$\alpha e^{\alpha(\tau-\mu)} - \frac{\beta}{\mu^2} e^{\beta(\lambda-\frac{1}{\mu})} = 0. \quad (5)$$

Note that when $\lambda \leq \frac{1}{\tau}$, it follows that $\mu^* = \tau$ since the provider's time constraint is not binding. In other words, when the arrival rate of patients is sufficiently low, providers can spend the time they need with each patient without incurring disutility from turning patients away or having patients leave voluntarily because of wait times. When $\lambda > \frac{1}{\tau}$, however, providers choose μ^* such that equation 5 holds.

We are centrally interested in how μ^* changes in response to an exogenous change in λ , which is captured in our empirical analysis by the reduction in clinic staffing to administer FluMist vaccines.

Using the implicit function theorem, we derive the following comparative static:

$$\frac{d\mu^*}{d\lambda} = \frac{\frac{\beta^2}{\mu^2} e^{(\beta\lambda - \beta/\mu)}}{\frac{\beta(2\mu - \beta)}{\mu^4} e^{(\beta\lambda - \beta/\mu)} - \alpha^2 e^{(\alpha\tau - \alpha\mu)}}, \quad (6)$$

such that the effect of a change in the arrival rate on the optimal amount of time spent with each patient is a function of preference parameters α and β , the current value of $(\tau - \mu)$, and the initial value of the arrival rate, λ .

This framework provides two key insights. First, the convex disutility of shortening visits equates to diminishing marginal returns with respect to average visit length. Providers are more willing to sacrifice time with patients when their average visit time is close to ideal than when it is considerably smaller. Second, conditional on a fixed μ , greater arrival rates will result in larger adjustments to μ^* ; however, this is somewhat misleading. As λ increases, we expect that providers will reduce μ^* , which will mute the effects of the increased arrival rate. In Figure 2, we therefore solve for μ^* for values of λ from 2 to 5 in 0.05 increments, and then present the first differences in μ^* as a numerical comparative static that takes into account changes in μ^* as λ increases. The observation that the numerical change in μ^* decreases as λ increases implies that increasing marginal disutility of $(\tau - \mu)$ dominates.

FIGURE 2 HERE

Figure 2 depicts how the comparative static of the optimal amount of time spent with patients changes under different conditions and different relative valuations of α and β . The two key takeaways here are: 1) $d\mu^*/d\lambda$ is negative and larger in magnitude when providers place more importance on seeing all patients relative to spending the “ideal” amount of time with each patient; and 2) exactly how providers will change μ^* in response to a FluMist induced change in λ will depend on the circumstances of the clinic in that day – including the arrival rate of patients.

This expository model therefore shows that exogenous reductions in capacity (i.e., an increase in the arrival rate of patients for remaining providers) can have different effects depending on the preferences of the provider and the volume of patients in the clinic that day. On relatively light days (i.e., when $\lambda < 1/\tau$), the clinic will have some amount of excess capacity. Since providers’

time constraints are not binding, there is no need to adjust the time they spend with each patient. On days when the clinic is closer to capacity, we expect a positive shock to the arrival rate to result in some decrease in μ^* . Finally, when the clinic is seeing very high numbers of patients, a change in λ is likely to have very little effect on μ^* as providers may be unwilling to sacrifice additional time with each patient.

5 Provider Responses to Decreased Capacity

Our initial results in Section 3 showed relatively little evidence of sufficient compensating behavior on behalf of the non-FluMist nurses. Our expository model suggests that this result could be driven by at least two factors: 1) there could be sufficient excess capacity already in the clinic such that clinics can absorb a temporary staff reduction without affecting actual patient care; or 2) nurses may exhibit a preference for time with patients over the number of patients seen. While several factors make a direct structural estimation of the provider’s optimization problem infeasible, we attempt to distinguish between these two explanations throughout this section.¹³

Our goal is to isolate situations in which capacity constraints are more likely binding and examine the effect of a reduction in capacity on such days. While we do not directly observe when constraints are binding, we attempt to identify such instances by exploiting variation in daily total visits to the clinic as well as exogenous short staffing. This approach arguably separates the role of nurse preferences from the role of built-in clinic capacity.

We pursue this approach with two additional models. First, at the provider level, we estimate an unconditional quantile regression with provider fixed effects to examine how the effect of FluMist on total number of patients seen varies over the distribution of patient volume (Firpo et al., 2009; Borgen, 2016). In this case, our fixed effects specification intuitively controls for time-invariant work characteristics of a given provider (i.e., nurse), and our quantile regressions investigate the different effects of FluMist days as the mean arrival rates also increases. We also include as

¹³One barrier in particular is that we do not observe people leaving the clinic. We instead only observe patients who ultimately received treatment at the clinic; although we do observe whether the visit was previously scheduled or was an unscheduled “walk-in” visit.

covariates a set of dummy variables for day of the week, year, month, and clinic, as in our initial estimates of equation (1). Estimates and 95% confidence intervals are presented in Figure 3.

FIGURE 3 HERE

For nurses removed from the clinic (the dashed line and respective confidence interval), we see no reduction in total patient volume for very low volume days. This is consistent with the notion that on sufficiently low demand days, a given provider may otherwise have some downtime. On these low demand days, being gone from the clinic for half the day does not substantially affect total visit volume, and administering FluMist vaccines essentially absorbs some of that downtime. As total patient volume increases, we see that being absent the clinic for half the day has a larger negative effect on the number of patients seen. In other words, on days that are busier than the median day, providers who are removed from the clinic see fewer than half the patients they otherwise would have. Of perhaps greater importance, among the nurses remaining in the clinic during FluMist days, we see no significant change in patient volume even on high volume days. While the magnitude of our estimates for the effect of FluMist on daily patient volume is larger on high volume days, the estimated effects are not large enough to offset the lost capacity from the FluMist nurses, and these estimates are never statistically significant for the nurses remaining in the clinic. Collectively, these results show that non-FluMist nurses do not fully compensate in the number of patients seen when staff is reduced on high volume days, suggesting that nurses prioritize time with individual patients over number of patients seen.

Second, we consider visit-level outcomes, where we model how the effect of FluMist on time with patients and other visit times change as we restrict the sample to increasingly high-volume days. Similar to our quantile regressions at the provider-level, this analysis focuses on days in which capacity constraints are more likely binding and offers additional insight on a provider's underlying preference for patients treated versus time with each patient. The differential effects of FluMist days by patient volume are presented graphically in Figure 4. The top panel presents the estimated effect and 95% confidence interval of FluMist on log number of minutes the nurse spends with a given patient, and the bottom panel presents results for log minutes of all other components of the

visit. Each line is constructed from a separate visit-level regression using ordinary least squares, analogous to that of equation (2), but where the estimation sample is limited only to those days with at least v visits in a day.

FIGURE 4 HERE

The results support a relationship between clinic capacity and visit length. While they are broadly consistent with the regression analysis in Table 3, they further inform that providers prioritize spending time with patients when capacity constraints bind. Specifically, as the number of visits per day increases, we initially see a small effect from a staffing reduction on nurse minutes but a substantial effect on other minutes (over 10% reduction). This effect on other minutes persists up to over 30 visits per day, or the 75th percentile of visit volume. Starting at 15 visits per day, providers’ time constraints begin to bind to where increased arrival rates from FluMist days reduces an RN’s time spent with patients. For days with total visit volume between 20 and 32 visits, RNs spend slightly over 5% less time with each patient, and even at the point where the estimated effect is largest (visit volume of 30), providers only reduce the time spent with patients by approximately 7% (or just over 2 minutes). However, consistent with our theoretical framework in Section 4, RNs do not further reduce time with patients on days where they are already sufficiently constrained (days with over 35 visits). This again suggests that providers strongly value time with patients over number of patients seen.

6 Robustness

While we contend that the administration of the FluMist vaccine was exogenous to any given clinic, it remains possible that other time-varying factors may be driving the selection of FluMist days from the KCHD central office. To examine this potential issue, we conducted placebo tests to verify that our results are driven by FluMist administration. To do so, we randomly draw 50 sets of placebo ‘FluMist’ days and compare our estimated coefficients in Section 3 to the distribution of estimated coefficients from the placebo ‘FluMist’ days.

Figure 5 presents the results. Effects of the true *FluDay* on total visit time and check-out time are greater than all placebo estimates (Figure 5a and Figure 5d).¹⁴ Looking at ready nurse time and nurse time, respectively, Figures 5b and 5c similarly show that over 95% of placebo coefficients have larger estimates than the true estimate. We conclude from these results that our estimates specifically from ‘FluMist’ days do not appear to be driven simply by random variation in visits over time but are instead reflective of some true underlying changes in nurse behaviors on FluMist days.

FIGURE 5 HERE

7 Discussion

In this paper, we exploit an exogenous source of variation in the capacity of public health clinics in the form of temporary staff reductions induced by FluMist days. Our results indicate that capacity reductions influence clinic behaviors along two margins: 1) on the extensive margin, clinics see fewer patients and prioritize scheduled visits over walk-ins; and 2) on the intensive margin, clinics first work to minimize administrative aspects of the visit but may ultimately reduce time with patients on high volume days. Overall, our findings indicate that providers value spending sufficient time with patients over seeing as many patients as possible.

In several aspects, we emphasize that these results represent a lower bound on the effect of capacity reductions, particularly when generalized to other areas of service provision. First, the service provided in the setting we study is fairly transactional (e.g., immunizations, disease screening, pregnancy tests, etc.). Most patients are referred to other providers if they have more nuanced or specialized needs. Because the nature of these visits is relatively simple within the health care context, there is less discussion/education to truncate than there may be in a family physician or hospital setting. For example, our results stand in contrast to Sharma et al. (2008), who examine provider behavior in emergency rooms. While emergency rooms are less able to delay care than public health clinics, they may be better able to adjust to increased demand by hastening dis-

¹⁴Since the estimates are negative, the true estimates are expected to be to the left end of the distributions.

charges. Second, our estimates only reflect the short run effects from temporary staff reductions. The nature of our exogenous variation does not capture longer-term compounding effects on the quality of care due to other factors such as provider fatigue from increased workload, absenteeism, or intention to quit.

Our results may offer some guidance as to the potential effects of staffing reductions in the provision of public services. Such reductions, even in the presence of some median-day excess capacity, are not without cost. We identify two responses to capacity reductions in particular. First, we find that providers maintain some minimum amount of time with customers such that remaining service providers do not fully compensate for the staffing reduction. We also find that providers prioritize scheduled visits over walk-in visits. The implication from these findings is that some customers go unseen. Second, while the reduction in time with customers is relatively small, the magnitude of reduction could be meaningful in certain settings.

Given our specific setting of public health clinics, each of these responses could carry important costs. For example, given that public health clinics immunize against communicable disease and treat sexually transmitted infections, untreated patients may generate substantial negative externalities. In addition, while a 5%-7% decrease in time with nurses may seem small, length of patients' time spent with providers has been shown to be a key determinant of 'quality of care' (Linzer et al., 2000; Whittington and McLaughlin, 2000; Wilson and Kaplan, 2000; Landau et al., 2007; Tai-Seale et al., 2007; Chen et al., 2009; Anand et al., 2011; Munyisia et al., 2011; McCloskey et al., 2014). For example, findings from Yarnall et al. (2003) suggest that a 5% reduction in time with patients would be sufficient to have otherwise counseled patients on STD prevention or contraception. Quantifying these responses in other contexts is a key piece of information if we are to understand the full effects of local, state, and federal budget decisions.

References

- Aiken, L., D. Sloane, J. Cimiotti, S. Clarke, L. Flynn, J. Seago, J. Spetz, and H. Smith (2010). Implications of the california nurse staffing mandate for other states: Nursing and home care. *Health Services Research* 45(4), 904–921.
- Anand, K., M. Paç, and S. Veeraraghavan (2011). Quality-speed conundrum: Trade-offs in customer-intensive services. *Management Science* 57(1), 40–56.
- Benz, M. (2005). Not for the profit, but for the satisfaction? evidence on workers well-being in non-profit firms. *KYKLOS* 58(2), 155–176.
- Borgen, N. T. (2016). Fixed effects in unconditional quantile regression. *Stata Journal* 16(2), 403–415(13).
- Bowblis, J. (2011). Staffing ratios and quality: An analysis of minimum direct care staffing requirements for nursing homes. *Health Services Research* 46(5), 1495–1516.
- Brueckner, J. K. (2002). Airport congestion when carriers have market power. *American Economic Review* 92(5), 1357–1375.
- Buurman, M., J. Delfgaauw, R. Dur, and S. VandenBossche (2012). Public Sector Employees: Risk averse and Altruistic? *Journal of Economic Behavior and Organization* 83, 279–291.
- Carlin, A. and R. E. Park (1970). Marginal cost pricing of airport runway capacity. *American Economic Review* 60(3), 310–319.
- Chalfin, A. and J. McCrary (2018). Are U.S. cities underpoiced? Theory and Evidence. *Review of Economics and Statistics* 100(1), 167–186.
- Chan, D. C. (2016). Teamwork and moral hazard: evidence from the emergency department. *Journal of Political Economy* 124(3), 734–770.
- Chan, D. C. (2018). The efficiency of slacking off: Evidence from the emergency department. *Econometrica* 86(3), 997–1030.

- Chen, L., W. Farwell, and A. Jha (2009). Primary care visit duration and quality: Does good care take longer? *Archives of Internal Medicine* 169(20), 1866–1872.
- Chen, M. and D. Grabowski (2015). Intended and unintended consequences of minimum staffing standards for nursing homes. *Health Economics (United Kingdom)* 24(7), 822–839.
- Cook, A., M. Gaynor, M. Stephens Jr, and L. Taylor (2012). The effect of a hospital nurse staffing mandate on patient health outcomes: Evidence from california’s minimum staffing regulation. *Journal of Health Economics* 31(2), 340–348.
- Daniel, J. I. (1995). Congestion pricing and capacity of large hub airports: A bottleneck model with stochastic queues. *Econometrica* 65(2), 327–370.
- DeAngelo, G. and B. Hansen (2014). Life and death in the fast lane: Police enforcement and traffic fatalities. *American Economic Journal: Economic Policy* 6(2), 231–257.
- Delfgaauw, J. and R. Dur (2008). Incentives and Workers’ Motivation in the Public Sector. *The Economic Journal* 118, 171–191.
- DeVany, A. (1976). Uncertainty, waiting time, and capacity utilization: A stochastic theory of product quality. *Journal of Political Economy* 84(3), 523–542.
- Dixit, A. (2002). Incentives and organizations in the public sector: an interpretive review. *Journal of Human Resources* 37, 696–727.
- Evans, W. and B. Kim (2006). Patient outcomes when hospitals experience a surge in admissions. *Journal of Health Economics* 25(2), 365–388.
- Evans, W. N. and E. G. Owens (2007). COPS and Crime. *Journal of Public Economics* 91(2), 181–201.
- Firpo, S., N. M. Fortin, and T. Lemieux (2009). Unconditional quantile regressions. *Econometrica* 77(3), 953–973.

- Freedman, S., T.-Y. Huang, D. Satin, and L. Smith (2018). Docs with their eyes on the clock? the effect of time pressures on primary care productivity. Working paper, Indiana University.
- Friedman, B. and M. V. Pauly (1981). Cost functions for a service firm with variable quality and stochastic demand: The case of hospitals. *The Review of Economics and Statistics* 63(2), 610–624.
- Ganster, D. and D. Dwyer (1995). The effects of understaffing on individual and group performance in professional and trade occupations. *Journal of Management* 21(2), 175–190.
- Gaynor, M. and G. F. Anderson (1995). Undertain demand, the structure of hospital costs, and the cost of empty hospital beds. *Journal of Health Economics* 14(1), 291–317.
- Gordon, T. and J. Iselin (2017, January). What everyone should know about their state’s budget. Technical report, Urban Institute, www.urbaninstitute.org.
- Hughes, D. and A. McGuire (2003). Stochastic demand, production responses and hospital costs. *Journal of Health Economics* 22(1), 999–1010.
- Kaiser Family Foundation (2013). Community health centers in an era of health reform: An overview and key challenges to health center growth. Technical report, the Henry J. Kaiser Foundation, <https://kaiserfamilyfoundation.files.wordpress.com/2013/03/8098-03.pdf>.
- Keeler, T. E. and J. S. Ying (1996). Hospitals costs and excess bed capacity: A statistical analysis. *The Review of Economics and Statistics* 78(3), 470–481.
- Landau, D.-A., Y. Bachner, K. Elishkewitz, L. Goldstein, and E. Barneboim (2007). Patients’ views on optimal visit length in primary care. *Journal of Medical Practice Management* 23(1), 12–15.
- Leete, L. (2001). Whither the nonprofit wage differential? estimates from the 1990 census. *Journal of Labor Economics* 19(1), 136–170.

- Lin, H. (2014). Revisiting the relationship between nurse staffing and quality of care in nursing homes: An instrumental variables approach. *Journal of Health Economics* 37(1), 13–24.
- Linzer, M., T. Konrad, J. Douglas, J. McMurray, D. Pathman, E. Williams, M. Schwartz, M. Gerity, W. Scheckler, J. Bigby, and E. Rhodes (2000). Managed care, time pressure, and physician job satisfaction: Results from the physician worklife study. *Journal of General Internal Medicine* 15(7), 441–450.
- Mani, V., S. Kesavan, and J. Swaminathan (2015). Estimating the impact of understaffing on sales and profitability in retail stores. *Production and Operations Management* 24(2), 201–218.
- Mas, A. (2006). Pay, reference points, and police performance. *Quarterly Journal of Economics* 121(3), 783–821.
- Matsudaira, J. (2014). Government regulation and the quality of healthcare: Evidence from minimum staffing legislation for nursing homes. *Journal of Human Resources* 49(1), 32–72.
- Mayer, C. and T. Sinai (MS2003). Network effects, congestion externalities, and air traffic delays: Or why not all delays are evil. *American Economic Review* 93(4), 1194–1221.
- McCloskey, R., C. Donovan, C. Stewart, and A. Donovan (2014). How registered nurses, licensed practical nurses and resident aides spend time in nursing homes: An observational study. *International Journal of Nursing Studies* 52(9), 1475–1483.
- Mello, S. (2019). More COPS, less crime . *Journal of Public Economics* 172(2), 174–200.
- Munyisia, E., P. Yu, and D. Hailey (2011). How nursing staff spend their time on activities in a nursing home: An observational study. *Journal of Advanced Nursing* 67(9), 1908–1917.
- Needleman, J., P. Buerhaus, V. S. Pankratz, C. L. Leibson, S. R. Stevens, and M. Harris (2011). Nurse staffing and inpatient hospital mortality. *New England Journal of Medicine* 364(11), 1037–1045.
- OECD (2017). *Government at a Glance*. Paris: OECD Publishing.

- Park, J. and S. Stearns (2009). Effects of state minimum staffing standards on nursing home staffing and quality of care. *Health Services Research* 44(1), 56–78.
- Prendergast, C. (2007). The motivation and bias of bureaucrats. *American Economic Review* 97(1), 180–202.
- Preston, A. E. (1989). the nonprofit worker in a for-profit world. *Journal of Labor Economics* 7(4), 438–463.
- Schilling, P. L., D. A. Campbell, M. J. Englesbe, and M. M. Davis (2010). A comparison of in-hospital mortality risk conferred by high hospital occupancy, differences in nurse staffing levels, weekend admission, and seasonal influenza. *Medical care* 48(3), 224–232.
- Sharma, R., M. Stano, and R. Gehring (2008). Short-term fluctuations in hospital demand: implications for admission, discharge, and discriminatory behavior. *RAND Journal of Economics* 39(2), 586–606.
- Tai-Seale, M., T. G. McGuire, and W. Zhang (2007). Time allocation in primary care office visits. *Health services research* 42(5), 1871–1894.
- Tong, P. (2011). The effects of california minimum nurse staffing laws on nurse labor and patient mortality in skilled nursing facilities. *Health Economics* 20(7), 802–816.
- Whittington, D. and C. McLaughlin (2000). Finding time for patients: An exploration of nurses’ time allocation in an acute psychiatric setting. *Journal of Psychiatric and Mental Health Nursing* 7(3), 259–268.
- Wilson, I. and S. Kaplan (2000). Physician-patient communication in hiv disease: The importance of patient, physician, and visit characteristics. *Journal of Acquired Immune Deficiency Syndromes* 25(5), 417–425.
- Yarnall, K. S. H., K. I. Pollak, T. Østbye, K. Krause, and J. L. Michener (2003). Primary care: Is there enough time for prevention? *American Journal of Public Health* 93(4), 635–341.

Yoe, J. (1988). The effects of workload and understaffing on staff in an educational service setting.
Journal of Environmental Psychology 8(2), 107–121.

Tables and Figures

Table 1: Summary Statistics for Clinic Visits^a

	Overall	FluMist Days	Non-FluMist Days
Total Clinic Days (N=1,713 with 101 FluMist observations)			
Total Visits	24.82 (11.03)	26.86 (9.85)	24.69 (11.09)
Scheduled Visits	8.07 (8.96)	10.06 (9.55)	7.96 (8.91)
Walk-in Visits	16.74 (10.68)	16.80 (11.74)	16.74 (10.62)
Components of Visit Length (N=42,514 with 2,713 FluMist observations)			
Total Visit Time	72.61 (49.99)	66.35 (43.22)	72.98 (50.34)
Check-in Time	11.25 (14.04)	10.23 (10.03)	11.31 (14.24)
Ready-nurse Time	10.65 (17.83)	9.76 (13.54)	10.71 (18.08)
Nurse Time	30.61 (28.70)	30.09 (24.71)	30.64 (28.95)
Ready-check-out Time	14.45 (28.62)	14.35 (28.90)	14.45 (28.61)
Visit/Patient Characteristics (%) (N=42,514 with 2,713 FluMist observations)			
Age Range			
0-10 yrs	18.32	20.32	18.19
11-20 yrs	20.72	17.18	20.96
21-30 yrs	26.53	26.40	26.54
31-40 yrs	15.70	15.86	15.69
41-50 yrs	7.96	7.89	7.97
51-60 yrs	5.45	6.23	5.39
61-70 yrs	3.45	3.87	3.43
71-80 yrs	1.44	1.55	1.44
81+	0.42	0.70	0.40
Reason for Visit ^b			
Immunization	33.48	37.82	33.19
STD Screen/Treat	16.79	22.12	16.42
Depo-Provera	5.95	4.83	6.03
Back-to-School Immunization	5.04	n/a	5.35
Travel Immunization	4.84	4.57	4.86
Day of Visit			
Monday	21.45	22.74	21.36
Tuesday	21.70	18.43	21.92
Wednesday	18.41	12.75	18.79
Thursday	19.30	22.67	19.07
Friday	19.15	23.41	18.86
Clinic Visited			
CDC	24.07	32.10	23.52
KCTE	9.96	7.96	10.10
KCWE	18.19	16.48	18.31
KCWH	22.35	16.00	22.78
TIC	25.43	27.46	25.29

^aStandard deviations in parenthesis.

^bLimited to the top 5 overall reasons observed in the data.

Table 2: Results for Provider-level Analysis^a

	FluNurse	FluDay ^b
Log Nurse Minutes	-0.439*** (0.061)	0.063 (0.046)
Log Total Visits	-0.446*** (0.049)	0.061* (0.033)
Log Walk-in Visits	-0.346*** (0.070)	-0.001 (0.037)
Log Scheduled Visits	-0.343*** (0.071)	0.095* (0.053)
Log Walk-in Share	0.113** (0.050)	-0.089** (0.038)

^aResults from a “within-estimator” with provider-level fixed effects. Column (1) reflects estimates for the coefficient on FluNurse based on the full sample. Column (2) presents estimates for the coefficient on FluDay when limiting the sample only to non-FluMist nurses. Different outcomes are presented in each row. Additional covariates excluded from the table include indicator variables for the clinic, day of the week, month of the year, and year. Standard errors in parenthesis clustered at the nurse level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

^bEstimates based on nurses who were in the clinic all day (i.e., excluding nurses who left the office to administer FluMist).

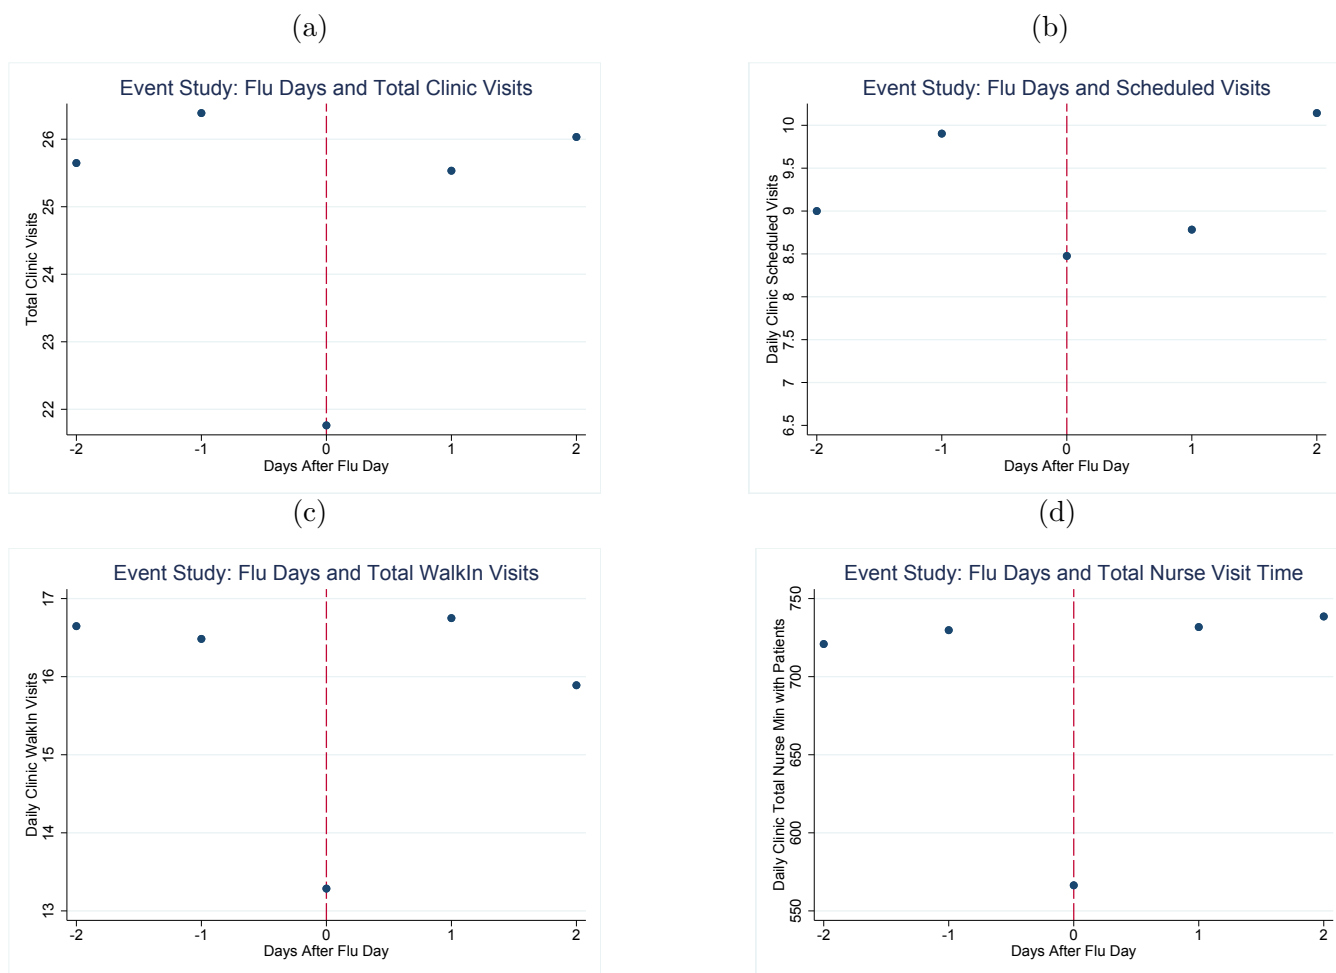
Table 3: Results for Visit-level Analysis^a

	All Visits	Non-FluMist Nurses ^b
Log Total Minutes	-0.071*** (0.018)	-0.077*** (0.020)
Log Check-in Minutes	-0.019 (0.048)	-0.016 (0.060)
Log Waiting Room Minutes	-0.058 (0.037)	-0.066 (0.042)
Log Nurse Minutes	-0.028* (0.016)	-0.016 (0.021)
Log Check-out Minutes	-0.105*** (0.034)	-0.081** (0.039)
Walk-in Visit	-0.105*** (0.030)	-0.110*** (0.035)

^aResults for the estimate on the FluDay coefficient based on ordinary least squares regressions. Column (1) reflects estimates from the full sample of all clinic visits, while column (2) presents results limited to non-FluMist nurses. Different outcomes are presented in each row. Additional covariates excluded from the table include indicator variables for the clinic, provider, reason for visit, age range of patient, day of the week, month of the year, and year. Standard errors in parenthesis clustered at the nurse level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

^bEstimates based on patients seen by nurses who were in the clinic all day (i.e., excluding nurses who left the office to administer FluMist).

Figure 1: Event Study: Daily Clinic Total Activity around Flu Mist Days



Each figure represents daily clinic totals of a given visit type or minutes of activity for all days within ± 2 business days from a day when nurses from that clinic went to administer Flu Mist.

Figure 2: First differences in μ^* as λ increases by 0.05




Figure 3: Quantile Regression Estimates on Log Total Visits by Total Visit Volume




Figure 4: Effects of FluMist on Length of Visit by Visit Volume


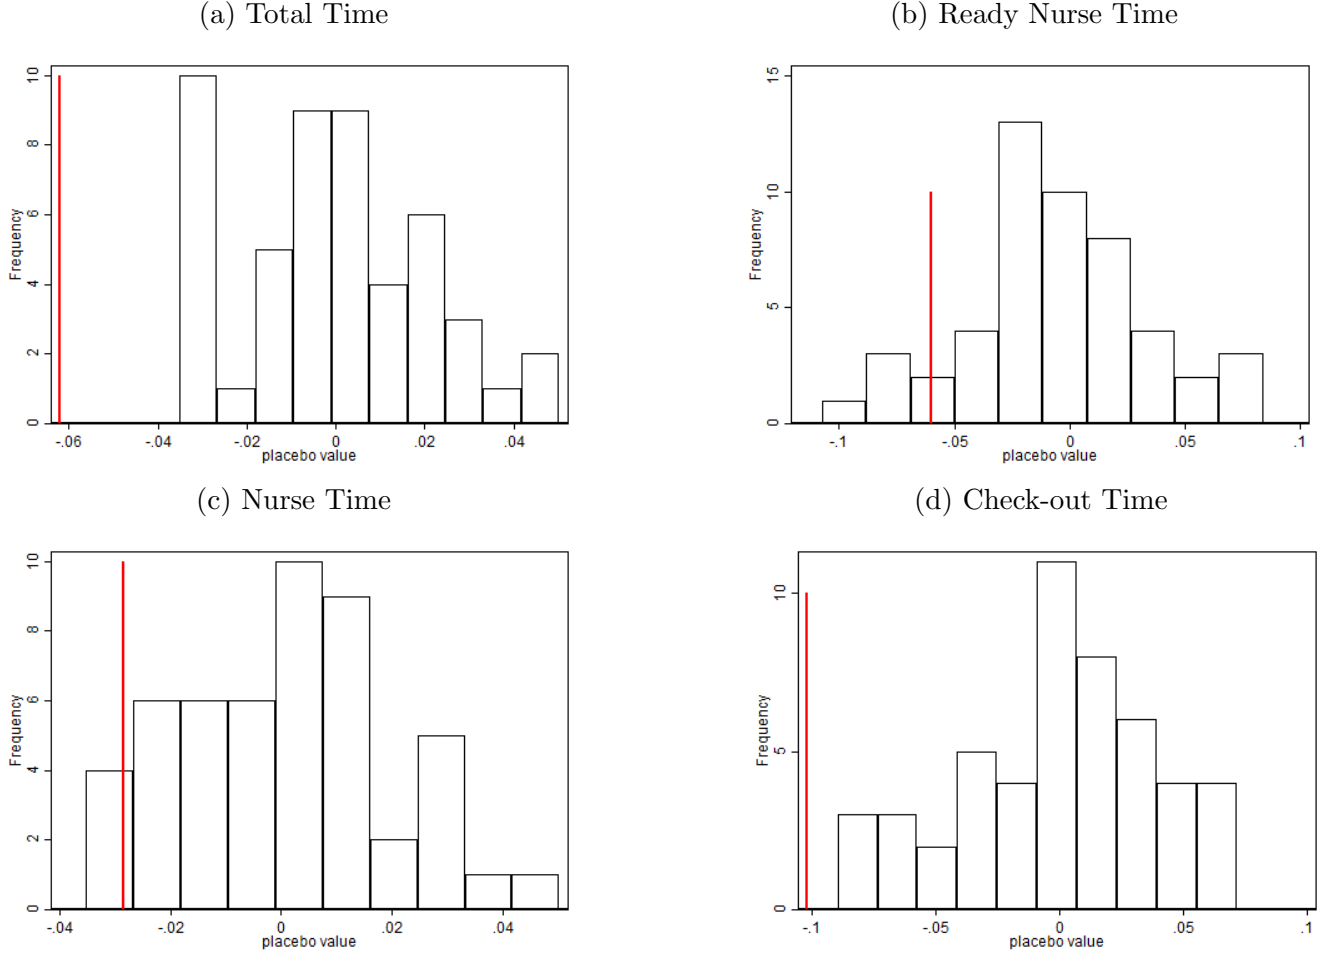


Figure 5: Placebo Tests



Each figure illustrates a distribution of β_t estimates from Equation 1 for the given outcome, where the distribution is generated by randomly sampling about 6% of the dates from the dataset. Solid lines represent the β_t estimate for *FluDay*, which can be found in column 1 of Table 3.