

NBER WORKING PAPER SERIES

ARE PUBLICLY INSURED CHILDREN LESS LIKELY TO BE ADMITTED TO HOSPITAL
THAN THE PRIVATELY INSURED (AND DOES IT MATTER)?

Diane Alexander
Janet Currie

Working Paper 22542
<http://www.nber.org/papers/w22542>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2016

We thank seminar participants at Yale, University of British Columbia, and Princeton for helpful comments, the staff at the New Jersey Department of Health and Senior Services for their assistance in accessing the data used in this paper, and the Center for Health and Wellbeing for financial support. All views expressed are those of the authors and cannot be attributed to the New Jersey Dept. of Health and Senior Services or any of its employees. 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Diane Alexander and Janet Currie. 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.

Are Publicly Insured Children Less Likely to be Admitted to Hospital than the Privately Insured
(and Does it Matter)?

Diane Alexander and Janet Currie

NBER Working Paper No. 22542

August 2016

JEL No. I13

ABSTRACT

There is continuing controversy about the extent to which publicly insured children are treated differently than privately insured children, and whether differences in treatment matter. We show that on average, hospitals are less likely to admit publicly insured children than privately insured children who present at the ER and the gap grows during high flu weeks, when hospital beds are in high demand. This pattern is present even after controlling for detailed diagnostic categories and hospital fixed effects, but does not appear to have any effect on measurable health outcomes such as repeat ER visits and future hospitalizations. Hence, our results raise the possibility that instead of too few publicly insured children being admitted during high flu weeks, there are too many publicly and privately insured children being admitted most of the time.

Diane Alexander
Federal Reserve Bank of Chicago
11th floor
230 S. LaSalle St.
Chicago, IL 60604
dalexander@frbchi.org

Janet Currie
Princeton University
316 Wallace Hall
Princeton, NJ 08544
and NBER
jcurrie@princeton.edu

Inequalities in income and inequalities in health are clearly associated: In particular, there is little doubt that poverty is associated with bad health outcomes. Yet, the extent to which the relationship is causal is controversial because causality could run in either direction and inference is often clouded by omitted variables. The existing literature is critiqued in Deaton (2001), Deaton and Lubotsky, 2003, and Deaton and Paxson (2001, 2004). These studies show that the relationship between income inequality and health inequality in aggregate data can disappear or even change sign depending on what else is included in the models. They make a strong case for careful investigations using individual-level data and research designs that can help to deal with potential confounding. This study continues this line of research by investigating the health effects of income-based inequalities in children's access to health care using unique longitudinal individual-level data.

In the U.S., most poor and near-poor children are covered by public health insurance under either Medicaid or SCHIP (State Child Health Insurance Program) while other children are covered by private health insurance policies offered by their parent's employers. Previous research has shown that for uncovered children, gaining access to public health insurance improved utilization of care and health outcomes (Currie and Gruber, 1996a,b; Dafny and Gruber, 2005; Selden and Hudson, 2006; Howell and Kenney, 2012; Meyer and Wherry, 2012; Brown et al., 2014). However, concerns about unequal access to care and disparities in health outcomes between those with public and private insurance are of long standing and remain unresolved.

It is difficult to tell whether patients with different types of insurance are being treated equally. For example, systematic differences in the unobserved characteristics of publicly and privately insured patients could drive any observed differences in treatment. Publicly insured

patients may also be needier, since on average they come from families with lower income and education, factors that are independently associated with poorer health. Moreover, publicly insured patients may be more likely to delay medical attention, and so may be in worse health when they do seek care. On the other hand, if publicly insured patients use hospital Emergency Rooms (ERs) as a source of routine care, then they may be less sick on average when they present at the ER (Gindi and Jones, 2014).

There may also be differences between publicly and privately insured patients in the extent to which they are likely to push back against the receipt of unnecessary care. There seems to be an unspoken assumption in much of the literature that the publicly insured should receive the *same* care as the privately insured. But it is possible that in many cases, the privately insured are over-treated and that less care is actually better. This possibility is especially relevant for hospitalization, which is costly, can be traumatic for both children and families, and exposes patients to the risk of hospital acquired infections and other hazards.

We investigate the issue of differential treatment of patients by insurance status using a unique data set based on the health records of all children between three months and 13 years of age who ever presented at a New Jersey hospital ER between 2006 and 2012. The first outcome we examine is whether or not a child was admitted to hospital. We rely on exogenous variation in the demand for hospital beds that is generated by local swings in the intensity of influenza (flu) outbreaks in order to separately identify the effects of hospital and patient behavior. When beds and other hospital resources are at a premium because of large influxes of adult flu patients, we find that publicly insured children who present at the Emergency Room (ER) are less likely to be admitted than in normal times. This result holds conditional on detailed diagnosis codes and within hospitals.

Our findings suggest that hospitals favor the privately insured for admission, especially when they are capacity constrained. However, rather than the poor being unjustly denied care when hospital beds are scarce, it could be that both privately and publicly insured children are receiving too many hospitalizations when beds are plentiful. We attempt to distinguish between these two scenarios by examining the effects of reductions in the probability of hospitalization during high flu weeks on subsequent ER visits and hospitalizations. We find that for the publicly insured, the marginal child who visits an ER during a high flu week is no more likely than a child who visits during normal times to reappear at the ER (or any other ER in New Jersey) within 14 days (or within 30 days). They are also no more likely to be admitted to hospital within the same time windows. These findings suggest that perhaps too many publicly insured children are admitted in normal times (and by extension, privately insured children who have even higher admission rates), rather than too few publicly-insured children being admitted in high flu weeks. If unnecessary hospitalizations are common, then they may be imposing substantial costs both on individual children and families, and on the public programs that pay for their care.

The rest of our paper proceeds as follows. We first discuss the background literature about the effects of public and private insurance. We then provide an overview of our data and research methods. Results are presented in section 5, and section 6 provides a discussion and conclusions.

2. Background

While some have argued that public insurance is worse than no insurance at all, there is little evidence to support this view (Gottlieb, 2011; Mocan, 2015). The uninsured fare worse on most measures of access to care (Institute of Medicine, 2001). However, the question of whether

publicly insured U.S. children on Medicaid or the State Child Health Insurance Program (SCHIP) have worse access to care than the privately insured continues to be controversial. Since Medicaid/SCHIP typically reimburses providers at lower rates than private health insurance plans, many doctors do not accept patients with public insurance or limit the number of these patients in their practices (Sloan, Mitchell, and Cromwell, 1978; Mitchell, 1991; Fossett et al., 1992; Newacheck et al., 1998; Decker, 2012). Nevertheless, a recent Department of Health and Human Services report concluded that privately and publicly insured children were similarly likely to have had a primary care visit in the past year (DHHS, 2012), a finding that is echoed in another recent study (Kenney and Coyer, 2012).

There may, however, be more subtle ways in which Medicaid/SCHIP children suffer worse access. For example, Kenney and Coyer (2012) found that publicly-insured patients were less likely to have a usual source of care with non-standard hours, which is one reason why these patients may be more likely to use emergency rooms for primary care (Garcia, 2010; Kangovi, 2013; Rhodes, 2013). A great deal of concern has also focused on inferior access to specialist care, which could lead to less preventive care and more medical emergencies (Bisgaier and Rhodes, 2011; Skinner and Mayer, 2007; U.S. GAO 2011; Wang et al., 2004).

Much less research has been devoted to the question of access to hospital care. Merrick et al. (2001) study children who were admitted to the hospital for asthma in California, Georgia, and Michigan. They find that publicly insured children received inpatient care of similar quality to privately insured children, conditional on being admitted. There is some evidence that children with public health insurance are more likely to be admitted to hospital than children without insurance (Dafny and Gruber, 2005). However, it is not known whether the publicly

insured are more or less likely to be admitted than the privately insured, other things being equal, or whether any differences in admission probabilities have effects on health outcomes.

Finally, there have been consistent concerns that many pediatric hospitalizations are unnecessary. Kemper (1988) conducted a retroactive assessment of visits to the University of Wisconsin Hospital and found that 21.4% of admissions were inappropriate. Shorter admissions were especially likely to be judged inappropriate (29%). Goodman et al. (1994) found that pediatric patients in zip codes with high per capita bed supply had 9% more admissions than children in areas with lower per capita bed supply. These findings suggest that many hospital admissions may not be beneficial to patients.

3. Data

The hospital discharge data come from the New Jersey Uniform Billing Records, which cover 2006-2012. Since our focus is on admission through the emergency room, our estimation sample consists of all records from general medical and surgical hospitals where the patient spent time in the emergency room (regardless of whether they were eventually admitted); that is, we include all records where there is an emergency room revenue code on the billing record. As we are particularly interested in how insurance type interacts with admission decisions for children, we further narrow our sample to visits where the patient is aged between three months and thirteen years. We exclude children under three months old, as the hospital experiences of infants (both in terms of realized experience and professional recommendations) is very different from that of older children.¹

¹ We expect very young children to not contribute much to our estimates, as congestion effects should be the most prominent for hospital visits that are marginal to admission. In the online appendix we exclude children less than one, and the results are nearly identical as in our main specification.

Figure 1 shows the age distribution of our sample, separately for publicly and privately insured children, for all ER visits (panel 1) and for admitted visits (panel 2). Figure 1 shows a very strong age-pattern for visits and admissions – younger children are more likely than older children to be admitted (this pattern is also reflected in admission rates in Figure 10). Also, after about two years of age, there are more privately insured than publicly insured children appearing in the ER (and being admitted). This pattern is due in part to a continuously increasing probability of being privately insured with age, as shown in Figure 2. However, even among the oldest children in our sample, over 30% are covered by public health insurance.

We are able to create a patient-level panel of hospital visits by matching patient records across visits by sex, date of birth, and first and last names. This process creates a unique patient identifier that allows us to follow each child over time and across hospitals. This feature is a particular strength of our data, as we can follow patients after their hospital visit, and see if they later returned to any hospital ER in New Jersey. This matching was performed on site in Trenton, and the data was then de-identified. The matching algorithm does a good job of catching slight misspellings, without lumping together people who appear to be separate individuals based on a manual inspection. More details on the matching process and the quality of the algorithm can be found in the online appendix.

The hospital discharge data include up to three detailed payer code variables for each visit (primary, secondary, and tertiary), which reflect the specific type of insurance used. Each payer code is classified into seven payer types: Medicare, Medicaid, Commercial, Blue Cross, Health Management Organizations (HMO), Self-Pay, and other. We re-categorize this information into two groups based on the primary payer: The first category, “public”, consists of Medicaid (Medicaid and Medicaid HMOs), NJ FamilyCare (New Jersey’s State Children’s

Health Insurance Program (SCHIP)), and the indigent. Indigent children make up just 1.16% of the public insurance category, and all results are robust to their exclusion. We include them in the public category because their expenses are likely to eventually be paid by public insurance even if they are not publicly insured at the time of admission. The “private” category consists of Blue Cross, non Medicaid/NJ FamilyCare HMOs, and other commercial insurers. Those with military insurance (CHAMPUS) are classified as “private” since their coverage is similar to that of the privately insured (CHAMPUS makes up 1.26% of the private insurance category). One strength of the data is that we can identify both Medicaid fee-for-service and Medicaid HMO patients as publicly insured patients. Further details about the construction of the categories and about Medicaid and SCHIP rules in New Jersey can be found in the online appendix.²

The hospital discharge data also include detailed diagnosis information for each visit, in the form of International Classification of Diseases, Clinical Modification (ICD-9-CM) diagnosis codes. In order to control flexibly for diagnosis (and for severity within broad diagnosis groups), we use the following grouping procedure. We start out at the most disaggregated level of the ICD-9-CM codes, and record both the number of visits and admission rates within each group. If there are more than 100 visits in the diagnosis group, and if the admission rate is between 5% and 95%, a dummy for the diagnosis code is included. If there are less than 100 visits, we aggregate up a level, and repeat the procedure, and continue the process until we are at the most aggregated level of diagnosis. Finally, we create categories that denote diagnoses with less than 5% admission rates, diagnoses with greater than 95% admission rates, and diagnoses with less

² The Affordable Care Act (ACA) is unlikely to affect our sample of children, as it primarily targeted adults, and the main components were implemented after 2012. That being said, the early portions of the ACA could indirectly affect the admission rates of children (e.g. through changes in hospital congestion). However, when we exclude 2010-2012 from our main sample, we find nearly identical results in the pre-reform period (online appendix).

than 100 visits when aggregated to the highest level. In the end, we are left with 415 diagnosis group dummy variables.

While children in our data are being treated for a wide range of conditions, the bulk of the visits for which admission is a real possibility are accounted for by relatively few diagnoses. Figures 3 and 4 list the top twenty most common diagnosis categories for all emergency room visits and admitted visits, respectively. Up to thirteen additional diagnoses may be recorded, but as most children just have one reported diagnosis, we control for the number of secondary diagnoses rather than including additional dummy variables for the secondary diagnosis categories. Figure 3 shows that most ER visits are for conditions with very low rates of hospitalization (such as a sprained ankle). The second panel focuses on diagnoses with higher rates of hospitalization and indicates that of these, the most common reasons for visiting the ER are “other;” unspecified noninfectious gastroenteritis and colitis (ICD-9-CM code 558.9); asthma not otherwise specified with acute exacerbation (ICD-9-CM code 493.91); pneumonia, organism not otherwise specified (ICD-9-CM code 486); urinary tract infection not otherwise specified (ICD-9-CM code 599.0); acute bronchiolitis due to other infectious organism (ICD-9-CM code 466.19); and dehydration (ICD-9-CM code 276.51).

Figure 4 shows that the top diagnostic categories for hospitalizations are similar to the top ER visit categories, with asthma, bronchiolitis, pneumonia, non-infectious gastroenteritis and dehydration again appearing prominently. Other important categories include acute appendicitis not otherwise specified (ICD-9-CM code 540.9) and bronchiolitis due to other infectious organisms (ICD-9-CM code 466.11).

One reason for privately insured children to have higher admission rates is that they are being treated for different conditions. However, Figure 5 shows that even conditional on particular diagnoses, privately insured children tend to have higher admission rates. Figure 5 ranks diagnoses by the admission rates of the publicly insured, and then plots the admission rates of the privately insured for each diagnosis for comparison. The size of the dots corresponds to the number of visits in the category. One can see an overall tendency for the dots representing privately insured patients with a particular diagnosis to lie above those for the publicly insured, an effect that is particularly pronounced for asthma admissions and in the tails of the distribution.

We construct weekly information about the intensity of influenza from the discharge data, using patients aged 18 and over. By using patients 18 and over, we avoid the possibility that children sick with the flu could crowd out other sick children with less severe conditions (leading to increases in child admission in some categories and decreases in others). Figure 6 shows the number of adult flu cases per week in the hospital data over our sample period, and Figure 7 shows flu intensity using the CDC “influenza-like-illness” measure for the mid-Atlantic region. Both figures are dominated by the several large spikes in the number of cases due to the H1N1 epidemic of 2009. The two time-series track each other quite well, making us confident that we are indeed able to pick up variation in the strength and intensity of flu seasons using the hospital discharge data.

The main advantage of constructing a measure of flu intensity from the hospital discharge data is that we can create local measures of influenza intensity based on the patient's zip code, rather than using regional data for the entire northeast. As we will show, there is considerable variation in the timing and intensity of flu shocks across different regions in New Jersey. For each zip code-week, the flu variable is the inverse-distance-weighted average of flu cases per

hospital bed at nearby hospitals. We define nearby hospitals as those within a ten-mile radius of each zip code centroid (though as the measure is inverse-distance-weighted, the 10 mile boundary ends up playing a very small role). Figure 8 shows the locations of New Jersey hospitals, as well as examples of the zip code-level hospital markets that are used to construct the flu intensity variable. One can see that hospitals are most plentiful in the heavily populated northeast of the state, and along the interstate-95 highway corridor from New York to Philadelphia.

Even within the small state of New Jersey, there is a lot of variation in the timing and intensity of flu outbreaks, especially relative to the number of available hospital beds. Figure 9 demonstrates this variation, by plotting the inverse-distance-weighted flu intensity measure over time in two different zip codes. Remarkably, there are almost twice as many flu cases per available hospital bed in Trenton during peak flu weeks as there are in New Brunswick, only 26 miles away. Moreover, in Trenton the last wave of the H1N1 flu season was the most severe, whereas in New Brunswick, the last wave was less severe than the first spike in H1N1 cases. It is possible that these differences reflect variations in vaccination rates in different populations in New Jersey, with African Americans and poorer people being less likely to be vaccinated than more affluent whites. Unfortunately, however, little reliable information about vaccination rates is available.

In Figure 10, we relate flu intensity to differential admission rates for publicly and privately insured children. The figure shows that privately insured children tend to have similar age-specific admission rates whether it is a high flu or a low flu week. However, publicly insured children over three years old are clearly less likely to be admitted in weeks when the local hospital flu burden is high.

One reason for publicly and privately insured children to be treated differently is that hospitals tend to receive larger payments from private insurers. While list charges are available for all observations, billing charges (the amount the hospital actually bills the insurer) are not well reported. In fact, most hospitals report the same number for the list charge and the billing charge. Billing charges are sensitive information; hospitals negotiate rates with insurers, and may not want insurers to know the amount other insurers are being charged for the same service. However, eight hospitals in New Jersey do report billing charges. These charges are consistently much lower than the list charges as one would expect given that insurers negotiate deep discounts from the list prices.

We use billing data from these eight hospitals to explore the differences between list and actual charges for Medicaid and privately insured patients, and to model the expected difference between list and actual charges in the other 65 New Jersey hospitals. We do this by estimating the following model using the eight hospitals that report billing data:

$$\begin{aligned}
 \text{Billing Chg}_{it} = & \beta_0 + \beta_1 \text{age}_{it} + \beta_2 \text{age}_{it}^2 + \beta_3 \text{any procedure}_{it} + \beta_4 \text{white}_{it} \\
 & + \beta_5 \text{black}_{it} + \beta_6 \text{hispanic}_{it} + \beta_7 \text{female}_{it} + \beta_8 \text{prev. respiratory visit}_{it} \\
 & + \beta_9 \text{prev. respiratory hospitalization}_{it} + \beta_{10} \text{prev. nonrespiratory visit}_{it} \\
 & + \beta_{11} \text{prev. nonrespiratory hospitalization}_{it} + \beta_{12} \text{list charge BC}_{it} \\
 & + \alpha_{\text{primary payer type}} + \alpha_{\text{diagnosis}} + \alpha_{\text{year}} + \alpha_{\text{month}} + \alpha_{\text{day of week}} + \varepsilon_{it}
 \end{aligned}$$

where *age* is age in days; *any procedure* is an indicator equal to one if the child received any procedure; the variables *female*, *white*, *black*, and *hispanic* denote characteristics of the child; the next four variables control for the child's previous medical history (in order to take account of

the fact that patient population varies between hospitals); *list charge BC* refers to the mean list charge for a child with no procedures and Blue Cross as the payer (in order to control for how expensive the hospital is) and the alphas denote vectors of fixed effects for payer type, diagnosis, year, month, and day of the week of the admission.

After estimating this regression, we use it to predict billing charges for the whole sample. This model does a good job of explaining the variation in billing charges for admitted patients, with an adjusted R-squared of 0.50. It has a harder time fitting the variation in billing charges for patients who are not admitted, and the adjusted R-squared drops to 0.11. For our purposes, however, being able to predict billing charges for admitted patients is the most important, as this is what determines incentives for admission. Note that all charges are expressed in real 2010 dollars.

In figure 11 we first show average *list* charges for publicly and privately insured patients. Within a hospital, all of the patients receiving the same procedure should have the same list charge. Figure 11 indicates that list charges are systematically higher for publicly insured patients than for privately insured patients at the same age. This gap may indicate that these patients are sicker on average, since they are receiving more services.

Figure 12 uses the predicted billing charges to explore differences in the payments hospitals are likely to receive for treating publicly and privately insured patients. It is clear that the private insurers are typically more generous, regardless of whether the patient is admitted or not. When the hospital has the capacity, then it will take both private and public patients. However, when beds are in short supply, hospitals have a clear reason to prefer privately insured patients.

Table 1 provides further information on patient characteristics by insurance type and flu intensity at the time of their ER visit. On average, 43 percent of children are covered by public insurance, and the average admission rate is 4.7 percent (Column 1). Columns 2 and 3 show the average differences between visits in high flu and low flu weeks (defining high/ low flu as zip code-weeks above/below the 90th percentile of zip code-weeks). Visits during high flu weeks are broadly similar to low flu weeks, though with more flu cases and lower admission rates for non-flu visits, which is consistent with some rationing of beds in high flu weeks. Columns 4 through 9 further disaggregate the sample by insurance type. Privately insured children are less racially and ethnically diverse than publicly insured children, and live in lower income zip codes. In addition, privately insured children have higher raw admission rates (column 5 versus column 8) — a disparity which rises during high flu weeks.

4. Empirical Specification

We have seen that the publicly insured are on average less likely to be admitted than the privately insured. In order to investigate this phenomenon more formally, we estimate the effect of insurance type on admission rates:

$$(1) \quad Admission_{it} = \beta_0 + \beta_1 I[public]_{it} + \beta X_{it} + \varepsilon_{it}$$

where $Admission_{it}$ is an indicator variable equal to one if the child is admitted, $I[public]_{it}$ is an indicator variable equal to 1 if the child is publicly insured, and the vector X_{it} includes dummies for age in months (which control very flexibly for changes in admission by age), dummy

variables for the detailed diagnosis codes, the number of diagnosis codes listed, year, month, and day of week; median income in the child's zip code (from the ACS); and the child's gender, race, and ethnicity (white, black, Hispanic). In some specifications, we additionally control for hospital fixed effects.

A limitation of (1) is that the publicly insured may face different admission rates than the privately insured because of differences in their unobserved characteristics rather than because of insurance status per se. Thus, we additionally examine whether a shock to the demand for hospital beds due to idiosyncratic spatial and temporal variation in flu severity exacerbates the disparity in admission rates. One can think of the variation in demand tracing out the supply of beds available to children as a function of the type of health insurance.

We implement this idea by estimating the effect of flu intensity, interacted with insurance type, on admission rates:

(2)

$$Admission_{izt} = \beta_0 + \beta_1 \mathbf{I}[\text{public}]_{it} + \beta_2 flu\ intensity_{zt} + \beta_3 (\mathbf{I}[\text{public}]_{it} * flu\ intensity_{zt}) + \beta X_{it} + \varepsilon_{izt}$$

where $flu\ intensity_{zt}$ measures the inverse-distance-weighted number of flu cases per hospital beds within 10 miles of a patient's zip code, and all other variables are defined as in equation 1. The parameter of interest in equation (2) is β_3 , which measures the differential effect of high flu intensity on the publicly insured. A negative value of β_3 indicates that publicly insured children are less likely to be admitted when demand for beds is high.

We estimate several versions of equations (1) and (2). We first look at children with flu diagnoses, and then at the full sample of all non-flu visits (visits where the primary diagnosis is not flu). We then use the same specifications to investigate the effect of flu intensity on total list charges (a measure of the value of the procedures the child actually received that should be unaffected by insurance coverage since the type of insurance normally determines discounts from the list price), the probability that a child returns to the emergency room within 14 or 30 days, and the probability a child returns to the emergency room and is admitted within 14 or 30 days.

5. Results

Table 2 shows our initial estimates of (1) and (2) first for ER visits with a flu diagnosis listed, and then for all other diagnoses. On average, 9.6% of children presenting with flu are hospitalized, but column 1 shows that publicly insured children are less likely to be admitted than the privately insured. Column 2 shows that even children with flu are less likely to be hospitalized in high flu weeks, suggesting that there are real constraints on the availability of beds at these times. The insignificant coefficient on the interaction term shows that there is however no significant differential effect of high flu weeks on the probability that publicly insured children with flu are admitted to hospital.

Column 3 of Table 2 shows that overall, about 4.7 percent of non-flu visits result in a hospital admission, and that publicly insured children are 0.24 percentage points less likely to be admitted than privately insured children with the same demographic characteristics, diagnoses, and time of visit. Overall, our equation (1) explains about 35% of the variation in admission probabilities for all non-flu diagnoses. Column 4 shows that including hospital fixed effects

actually raises the admission penalty associated with public insurance, which suggests that there is more difference within hospitals than across hospitals.

Columns 5 and 6 show estimates of equation (2) which add the controls for the number of flu cases per bed locally, and the interaction of this variable with the indicator for public health insurance. The main effect of public insurance remains much the same as in the corresponding specifications in columns 3 and 4. The coefficient on flu cases per bed shows that, as expected, children are generally less likely to be admitted in high flu seasons. Table 1 shows that the mean number of flu cases per bed in the top decile of flu weeks is 0.093. Multiplying this number by the coefficient of -0.0111 on flu cases per bed in column 5 suggests that the probability of being admitted declines by 0.001 in high flu weeks on a baseline of 0.047, a relatively small effect. However, for the publicly insured, the decline in admissions is three times larger in high flu weeks at 0.003, reducing admission probabilities from 4.7 to 4.4 percent. This represents a reduction of about 9,700 child hospital admissions among the publicly insured. The coefficients in column 5, with hospital fixed effects, also suggest a public insurance penalty on admissions, which grows larger in high flu weeks.

b) Effects on List Charges

Table 3 examines the effect of public insurance on total list charges. We use total list charges because list charges are generally independent of insurance status, and so any difference reflects differences in the quantities and types of services received rather differences in price per service actually charged. Column 1 shows that the list charges per ER visit are \$107.85 cheaper for publicly insured children, which is consistent with their lower probability of being admitted even within hospital. Column 2 shows that list charges fall even more steeply in high flu weeks.

At the “high flu” mean cases per bed of 0.093, the coefficient of -560.86 translates into charges that are an additional \$52.15 lower per visit. The coefficient on the interaction between public insurance and flu cases per bed indicates that list charges are even lower (by \$62.22) for the publicly insured in high flu weeks. Thus, column 2 suggests that within hospitals, publicly insured children run up lower list charges than the privately insured, consistent with being less likely to be admitted, and that this differential is greater in high flu weeks.

c) Effects on Future ER Visits and Hospitalizations

The results discussed so far raise an important question: Are the publicly insured receiving too few hospitalizations? It is at least possible that the privately insured are receiving too many hospitalizations. It may be the case that public insurers and parents of publicly insured children are less likely to decline unnecessary hospitalizations. Clearly these alternative scenarios have very different implications for public policy.

One way to try to distinguish between these stories is to examine health outcomes. If children who need hospitalization are being turned away, then we should expect them to return to the ER and/or to be hospitalized soon after. Tables 4 and 5 examine these outcomes, where the variable “return to ER” is equal to one if a child shows up back in a New Jersey emergency room soon after the initial visit, and future hospital admissions are similarly defined as a child being admitted in a New Jersey hospital within a short window after the initial visit. The first row of Table 4 shows that publicly insured children who present in the ER are more likely than privately insured children to return to the ER. As the window of observation is extended from 14 to 30 days (comparing columns 1 and 2 to columns 3 and 4) this effect grows. This main effect may reflect the subset of publicly insured children who tend to use ERs as a source of primary care.

The second row of Table 4 shows that in weeks with many flu cases per bed, children are more likely to return to the ER within 14 days, though less likely to return within 30 days. The 14 day estimate suggests that hospitals are genuinely capacity constrained during high flu weeks and that some children who are not admitted when they first present at the ER need hospital care. The negative effect at 30 days may reflect seasonality in flu in that high flu weeks may be followed by low flu weeks in later months (as in Figures 6, 7, and 9).

Turning to the key interaction term between flu cases per bed and public health insurance, the estimated coefficient is not statistically significant. This means that although we have established above that publicly insured kids are less likely to be hospitalized in high flu weeks, this reduction in hospitalization apparently has no effect on the probability of returning to the ER at either 14 days or 30 days following the initial visit.

Similarly, Table 5 shows that although there is a very small positive main effect of public insurance on future hospital admissions within 14 days or 30 days of the index visit, there is no significant of the interaction of flu cases per bed and public insurance on future admission probabilities.

Together these null results on future ER visits and admissions suggest that the reductions in hospitalizations among the publicly insured during high flu seasons have no adverse effect on their health. In turn, this finding suggests that a real problem for the publicly insured, at least in terms of their treatment at hospitals, may be that they are too likely to be admitted when hospitals have excess capacity rather than being too likely not to be hospitalized during high flu seasons when beds are in high demand.

d) Robustness checks

Our preferred measures of a child returning to the ER or to be admitted do not require that the return visit be in the same diagnosis group as the original visit. We use this broad measure rather than requiring the return visit to be within the same ICD-9-CM code because the codes are quite narrow, and complications that may arise from not being admitted would not necessarily be recorded in the same code as the original visit. The downside to using this broad measure is that we are likely picking up some noise from children who return to the hospital for reasons completely unrelated to the index visit---for example, if they were involved in a car accident. In the online appendix we explore this issue, first by excluding return visits with injury codes as the primary diagnosis (ICD-9-CM 800-957), and second by looking only at visits within the broad category of diseases of the respiratory system (ICD-9-CM codes 460-519). In both cases, we continue to find null results, suggesting that any bias introduced into our main estimates by unrelated hospital visits following in quick succession is minimal.

Finally, while our main estimates report the average effect across all patient types, there is also scope for heterogeneity, especially with respect to types of diagnoses. In the online appendix, we rerun our analysis separately on visits for diagnoses with mid-range admission rates (5-95%) and low admission rates (less than 5%). Since doctors treating patients in diagnoses with mid-range admission rates face a more real choice over admission, we expect to find larger negative effects on the likelihood of admission for publicly insured children with mid-range admission rates during high flu times---and indeed we find a nearly identical pattern of results, with point estimates 15-30% larger than in the full sample. However, we still find no effect of these lower admission rates on health outcomes, as measured by the probability that the children return to the ER or are admitted to the hospital soon after the initial visit. Conversely,

we find negative but small and statistically insignificant congestion effects for publicly insured children with low admission rate diagnoses, which makes sense as very few of these children are likely to be on the margin of admission.

6. Discussion and Conclusions:

There is continuing controversy about the extent to which publicly insured children are treated differently than privately insured children, and whether differences in treatment matter. We show that on average, hospitals are less likely to admit publicly insured children than privately insured children who present at the ER. This pattern is present even after controlling for detailed diagnostic categories, suggesting that it is not simply a matter of publicly insured patients being less sick on average. Controlling for hospital fixed effects increases the size of the public insurance “penalty”, suggesting that it cannot be explained by differences between the hospitals used by the privately and publicly insured. Another possibility is that the privately insured are simply more profitable than the publicly insured, so that providers prefer to admit privately insured patients, other things being equal. A related hypothesis is that the privately insured (or their insurers) are more likely to accept unnecessary hospitalizations. Since unnecessary hospitalizations are costly both to society and to families, and may expose children to unnecessary risks, it is important to distinguish between these possibilities.

We explore these questions using variation provided by influenza epidemics. The flu increases demand for hospital beds in a way that varies across both time and space. And given that we focus on adult influenza cases, these demands are arguably exogenous to other factors that affect children’s probability of being hospitalized for conditions other than flu.

We show that the difference between the probability that privately insured children are hospitalized and the probability that publicly insured children are hospitalized is larger in high

flu weeks. This result suggests that hospitals prefer to fill available beds with privately insured children, perhaps because they are more lucrative. One can think of the available beds being allocated first to the privately insured, with any remaining beds being given to the publicly insured. However, by itself this result cannot tell us whether it is good or bad for publicly insured children to have a lower probability of hospitalization in high flu weeks.

If publicly insured children were suffering from the reduction in admission probabilities during high flu weeks, then we might expect to see evidence of harmful effects in the form of returns to the ER and future hospital admissions. However, we see no such patterns. Hence, our results raise the possibility that instead of too few publicly insured children being admitted during high flu weeks, there are too many publicly and privately insured children being admitted most of the time.

References

Bisgaier, Joanna and Karin V. Rhodes, "Auditing Access to Specialty Care for Children with Public Insurance," New England Journal of Medicine, June 16, 2011, 2324–2333, <http://www.nejm.org/doi/full/10.1056/NEJMsa1013285>

Currie, Janet and Jonathan Gruber. "Saving Babies: The Efficacy and Cost of Recent Expansions of Medicaid Eligibility for Pregnant Women," The Journal of Political Economy, December, 1996, 104 #6, 1263-1296.

Currie, Janet and Jonathan Gruber. "Health Insurance Eligibility, Utilization of Medical Care, and Child Health," The Quarterly Journal of Economics, May 1996, 111 #2, 431-466.

Currie, Janet and Duncan Thomas "Medical Care for Children: Public Insurance, Private Insurance, and Racial Differences in Utilization," The Journal of Human Resources, Winter, 1995, 30 #1, 135-162.

Dafny, Leemore and Gruber, Jonathan. "Public Insurance and Child Hospitalization: Access and Efficiency Effects," Journal of Public Economics 89 (2005) 109– 129.

Deaton, A. "Inequalities in Income and Inequalities in Health." In The Causes and Consequences of Increasing Inequality. Chicago: University of Chicago Press (2001).

Deaton, A, and D Lubotsky. "Mortality, Inequality and Race in American Cities and States." Social Science and Medicine 56 #6 (2003): 1139-1153.

Deaton, A, and C Paxson. "Mortality, Education, Income and Inequality Among American Cohorts." In Themes in the Economics of Aging. Chicago : University of Chicago Press (2001).

Deaton, A, and C Paxson. "Mortality, Income, and Income Inequality over Time in Britain and the United States." In Perspectives on the Economics of Aging. Chicago : University of Chicago Press (2004).

Decker, S. "In 2011, Nearly One-Third of Physicians Said They Would Not Accept New Medicaid Patients, But Rising Fees May Help," Health Affairs, v. 31 #8, August 2012, 1673–167

Dubay, L and G. Kenney, "Health Care Access and Use among Low-Income Children: Who Fares Best?" Health Affairs 20 #1, 2001.

Fossett, J.W. et al., "Medicaid and Access to Child Health Care in Chicago," Journal of Health Politics, Policy and Law, 17 #2, 1992, 273–298.

Freeman and Corey, "Insurance Status and Access." Health Services Research, Dec 1993; 28(5): 531–541.

Garcia T et al., “Emergency Department Visitors and Visits: Who Used the Emergency Room in 2007?” NCHS Data Brief No. 38, May 2010.

Gindi R. and L. Jones, “Reasons for Emergency Room Use Among U.S. Children: National Health Interview Survey, 2012,” NCHS Data Brief, No. 160, July 2014

Goodman, David, Elliott Fisher, Alan Gittelsohn, Chiang-Hua Chang, Craig Fleming. “Why are Children Hospitalized? The Role of Non-Clinical Factors in Pediatric Hospitalizations,” Pediatrics, June, v93 #6.

Gottlieb S, “Medicaid is Worse than No Coverage at All,” Wall Street Journal online, March 10, 2011. <http://online.wsj.com/article/SB10001424052748704758904576188280858303612.html>

Howell, E and G. Kenney, “The Impact of the Medicaid/CHIP Expansions on Children: A Synthesis of the Evidence,” Medical Care Research and Review 69 #4, August 2012.

Institute of Medicine, Coverage Matters: Insurance and Health Care, Committee on the Consequences of Uninsurance, Board on Health Care Services, Washington D.C.: National Academy Press, 2001.

Kemper, Kathi. “Medically Inappropriate Hospital Use in a Pediatric Population,” New England Journal of Medicine, 1988, April 21, 1033-1037.

Smedley, Brian D., Adrienne Y. Stith, and Alan R. Nelson, Editors, Unequal Treatment: Confronting Racial and Ethnic Disparities in Health Care, (Washington D.C.: Institute of Medicine Committee on Understanding and Eliminating Racial and Ethnic Disparities in Health Care) 2003.

Kaestner, R. “Medicaid and Children’s Access to Care,” Journal of the American Medical Association, 281 #14, 1999, 1273.

Kangovi, S. et al. “Understanding Why Patients of Low Socioeconomic Status Prefer Hospitals over Ambulatory Care,” Health Affairs, 32 #7, July 2013.

Kenney, G. and C. Coyer, “National Findings on Access to Health Care and Service Use for Children Enrolled in Medicaid or CHIP, MACPAC Contractor Report #1, March 2012.

Merrick, Nancy, Robert Houchens, Sandra Tillisch, and Bruce Berlow, “Quality of Hospital Care of Children with Asthma: Medicaid Versus Privately Insured Patients,” Journal of Health Care for the Poor and Underserved, v. 12 #2, 2001, pp. 192–207

Mitchell, J. “Participation in Medicaid Revisited,” Medical Care, July 1991, 645–653.

Mocan, N. et al. “The impact of mothers’ earnings on health inputs and infant health,” Economics and Human Biology, 19, 2015, pp. 204-223.

Newacheck, P. et al., "The Role of Medicaid in Ensuring Children's Access to Care," Journal of the American Medical Association 280, no. 20 (1998): 1789–1793.

Rhodes, K. et al. "Patients Who Can't Get an Appointment Go to the ER': Access to Specialty Care for Publicly Insured Children," Annals of Emergency Medicine, 61 #4, April 2013.

Rosenbach, M.L. "The Impact of Medicaid on Physician Use by Low-Income Children," American Journal of Public Health 79 #9, 1989, 1220–1226.

Selden, T. and J. Hudson, "Access to Care and Utilization Among Children: Estimating the Effects of Public and Private Coverage," Medical Care, 44 #5, 2006.

Short, P.F. and D.C. Lefkowitz, "Encouraging Preventive Services for Low-Income Children: The Effect of Expanding Medicaid," Medical Care 30 #9, 1992, 766–780.

Sloan, F., R. Mitchell, and J. Cromwell, "Physician Participation in State Medicaid Programs," Journal of Human Resources, Supplement 1978, 211–245.

Skinner, A. and M. Mayer, "Effects of Insurance Status on Children's Access to Specialty Care: A Systematic Review of the Literature," BMC Health Services Research, 7 #194, 2007.

U.S. Dept. of Health and Human Services, 2012 Annual Report on the Quality of Care for Children in Medicaid and CHIP, , Dec. 2012.

U.S. General Accounting Office, "Medicaid and CHIP: Most Physicians Serve Covered Children but have Difficulty Referring them for Specialty Care," June 2011.

Wang, Edward, Meeryo Choe, John Meara, and Jeffrey Koempel, "Inequality of Access to Surgical Specialty Health Care: Why Children with Government-Funded Insurance Have Less Access than Those with Private Insurance in Southern California," Pediatrics, v. 114 #5, 2004, e584–e590.5.

Weissman, J. et al. "The Quality of Hospital Care for Medicaid and Private Pay Patients," Medical Care, 51 #5, May 2013.

Appendix:

New Jersey Medicaid and CHIP Rules for Children

From state of NJ website and “Holding Steady, Looking Ahead: Annual Findings of a 50 State Survey of Eligibility Rules, Enrollment and Renewal Procedures, and Cost Sharing Practices in Medicaid and CHIP 2010-2011”

Eligibility

- The upper income limit of New Jersey’s combined Medicaid/SCHIP program is 350% of the FPL.
- Income eligibility limits for Medicaid are 200% FPL for children 0-1, 133% FPL for children 1-19. Income eligibility limit for CHIP is 350% FPL for children ages 0-19. The separate CHIP program is for children not eligible for Medicaid
 - Infants born to Mothers enrolled in Medicaid are covered up to 200% FPL; infants born to non-Medicaid covered mothers are covered up to 185% FPL
 - Families above the 350% cut off for CHIP can buy into the program for a monthly premium of \$144. The child must have been uninsured for 6 months prior to enrolling
- Pregnant women eligible for Medicaid at 185% FPL and CHIP at 200% FPL. An asset test is not required, and the program operates under presumptive eligibility. Pregnant immigrants are covered regardless of immigration status.
- There is the same eligibility system for Medicaid and CHIP

Application

- Streamlined application: there is a joint Medicaid/CHIP application, which requires no face-to-face interview or asset test. Applications are available online, and can be submitted electronically.

Enrollment

- Before enrollment in CHIP, a child must be uninsured for 3 months
- Streamlined enrollment: Medicaid and CHIP both operate on presumptive eligibility and use express lane eligibility for Medicaid and CHIP.
 - New Jersey implemented Express Lane Eligibility for Medicaid 5/1/09, and for CHIP 9/2012: “In an effort to avoid requiring families to provide the same information to multiple programs and to achieve administrative efficiencies, ELE allows states to use income and other eligibility findings from another assistance program as evidence of eligibility for Medicaid and CHIP. New Jersey uses data

from NJ 1040 tax form.”

Renewal

- Frequency of renewal for Medicaid/CHIP: 12 mo, no face to face interview required
 - There is a joint Medicaid/CHIP renewal form; Medicaid renewal uses the express lane since 5/1/09; CHIP in 2012
- Both Medicaid and CHIP also have 12-month continuous eligibility
- Disenrollment policies: grace period for non-payment is 60 days, with no lock-out period. To reenroll, must reapply for coverage and repay outstanding premiums.

Premiums and Copays

- Premiums in CHIP start at 201% FPL; copays start at 151% FPL
 - At 201% FPL premiums are about \$40/mo, at 250% FPL \$79/mo, at 301% FPL \$133/mo
- Copays for services: At 151% FPL there is a \$5 copay for non-preventive physician visits, \$10 copay for ER visit, 0\$ copay for inpatient. At 201% FPL the physician visit and inpatient visit copays stay the same, and ER goes up to \$35
- Copays for drugs: At 151% \$1 for generic, \$5 for brand name; at 201% \$5 for all.

Information on Coding Medicaid/NJ FamilyCare HMOs:

We categorize the payer information into two groups based on the primary payer: Medicaid, NJ FamilyCare, indigent, and other government (Medicaid, Medicaid HMOs, FamilyCare HMOs, Title XIX Medicaid, other government, and indigent which comes from the “other” category), and private (Commercial, Blue Cross, non Medicaid/FamilyCare HMOs, Champus, and New Jersey State Health Benefits).

The HMO category is broken into Medicaid/FamilyCare and non-Medicaid/FamilyCare HMOs based on information about product lines from New Jersey HMO contracts. Six HMOs with Medicaid/NJ FamilyCare product lines were identified to be in operation during the time period. Four of these HMOs had no commercial product lines, and were easily classified as Medicaid/FamilyCare HMOs. The other two have both Medicaid/FamilyCare and Commercial product lines, and in the data there is no way to distinguish which patients are Medicaid/FamilyCare and which are not. All patients with these HMOs as primary payers were coded as Medicaid/FamilyCare, though some are likely private. The results are robust to whether these patients are coded as Medicaid/FamilyCare or private.

HMO Contracts with Medicaid Product Lines	2006	2007	2008	2009	2010	2011	2012	Commercial Product Line
AmeriChoice of New Jersey/UnitedHealthcare	x	x	x	x	x	X	x	
AMERIGROUP New Jersey/Americaid	x	x	x	x	x	X	x	
Healthfirst Health Plan of New Jersey, Inc				x	x	X	x	
Health Net of New Jersey, Inc.	x	x	x	x				x
Horizon Healthcare of New Jersey/ HMO Blue	x	x	x	x	x	X	x	x
University Health Plans, Inc.	x	x	x	x				

Information on HMO contracts and product lines from the 2006-2012 New Jersey HMO Performance Reports (Report Cards)

<http://www.state.nj.us/dobi/lifehealthactuarial/hmo2007/index.html>

Information on the String Matching Algorithm

The matching algorithm creates a patient identifier by finding records with the same date of birth and the same or very similar first and last names. Specifically, the Levenshtein edit distance is used to match names, because of problems with typos and misspellings (stata command `strgroup`). While it is possible that we are picking up a few cases of different people with the same name and birthday, it does not seem to be a large problem. To examine the quality of the matching algorithm, we looked all visits of children included in the main sample, during the age window of 275 days around the first birthday.

The main worry is that there may be many children with similar names and the same birthday, who are being aggregated by the algorithm. In order to assess whether this was a concern, we looked at people with the most common first and last name combinations. We took first and last name combinations that the algorithm assigned to at least eight children, and called this the sample of “common names”.

In order to assess the match quality, we looked at the three-digit zip code of residence reported for each visit. Of this sample of children with extremely common names, those with more than ten visits reported all visits in the same three-digit zip code. This suggests that the algorithm did not mistakenly aggregate people with common first and last name combinations – otherwise we would expect the people with common name combinations and many visits to report multiple zip codes. Of all people with these common names, 92.49% reported just one three-digit zip, 6.94% reported two, and 0.58% reported three.

Furthermore, we manually inspected all patients in the top 1% of number of visits. None of these children had “common names”, as defined above, and almost all reported either just one three-digit zip code, or a combination of neighboring three digit zip codes.