The Vaccination Kuznets Curve: Rise and Fall of Vaccination Rates with Income

Yutaro SAKAI[*,†]

This version: October 30, 2016

Job Market Paper

Abstract

This paper presents a new stylized fact about the relationship between income and childhood vaccination. It shows vaccination rates first rise but then fall as income increases. This pattern is observed in WHO country-level panel data, and in US county-level panel and individual-level repeated cross-section data. This pattern suggests that both low and high-income parents are less likely to follow the standard vaccination schedule, although for different reasons. To shed light on these parents' vaccination decisions, I develop a simple model and show how substitutes for vaccination such as avoidance measures and medical care could be responsible for this finding.

JEL: H41, I12, I15
Keywords. avoidance, income, immunization, infectious disease, medical care, NIS, childhood vaccination

[*]PhD candidate, Department of Economics, University of Calgary, 2500 University Dr. NW, Calgary, Alberta, Canada, T2N 1N4, ysakai@ucalgary.ca

[†]This paper was previously titled "Who Makes Us Sick and Why: Declining Childhood Vaccination Rates in High-Income Communities." I am grateful for my supervisor, Dr. M. Scott Taylor, for his continuous support and excellent guidance. The paper has also benefited from comments by Christopher Auld, Daniel Gordon, Arik Levinson, Arvind Magesan and seminar participants at the University of Calgary.

1 Introduction

Vaccination is one of the greatest inventions in human history. This medical intervention is so powerful that during the 1940s-1950s, people believed that the war against infectious diseases was almost over.[1] Since then, however, we have failed to eradicate any disease—except for smallpox—and the outbreak of various diseases continues to occur in even the most developed countries. Effective vaccines are indeed available for many diseases, but vaccination rates are not high enough to prevent outbreaks. Moreover, there is a growing concern that vaccination rates are actually falling in some segments of the population, particularly, in high-income groups.

This paper examines the relationship between childhood vaccination rates and per capita income using three datasets that differ in their level of aggregation and country coverage. At the country and county level, I examine the relationship between childhood vaccination rates and per capita income of the region. At the individual level, I examine the relationship between the probability of a child being up-to-date with a vaccine schedule and per capita income of the family. While the raw data shows vaccination rates typically increase with per capita income, I show that vaccinate rates initially rise but then begin to fall as per capita income increases conditional on region and year fixed effects.[2] This consistent pattern across datasets indicates that both low- and high-income parents are less likely to follow the standard vaccination schedule.

The inclusion of region- and year-fixed effects in the estimation is therefore key to these results. Region-fixed effects capture any region-specific factors that are constant within each region over time. For example, good local institutions may help increase both income and vaccination rates, creating a spurious positive correlation between these variables in the raw data. Region-fixed effects will help eliminate this correlation. Similarly, year-fixed effects allow for secular changes in vaccination rates over time that are common to regions. For example, the knowledge about the importance of vaccines changes over time while simultaneously regions become richer. This secular change can again create a spurious positive correlation in the raw data. As a result, a method that allows for both region- and year-fixed effects may uncover quite different relationships between per capita income and vaccination rates from what the raw data suggests.

To investigate this possibility, I employ three different datasets. The country-level data is obtained from the World Health Organization (WHO) that consists of 70 countries with six vaccines for the 1980-2014 period. The county- and individual-level datasets are obtained from the United States (U.S.). The county-level dataset includes the period 1995-2008, consisting of 229 counties with seven vaccines. The individual-level dataset is provided by the National Immunization Survey (NIS), which includes the vaccination status of eight childhood

[1]In 1948, U.S. Secretary of State George Marshall expressed his view that the conquest of all infectious diseases were imminent. This great optimism was due not only to vaccination but also to the growing number of antibiotics and the discovery of other chemicals that could effectively kill mosquitoes and other insect pests (Garrett, 1994).

[2]In the country-, county-, and individual-level analysis, the term region corresponds to a country, county, and state, respectively.

vaccines for >160,000 individuals during the 2005-2014 period. While the results naturally differ across these quite different samples, vaccination rates and per capita income often exhibit a hump-shaped relationship.[3] Vaccination rates peak within the range of $25,000-$40,000 in most cases.

But why do parents in both tails of the income distribution not vaccinate their children? Low-income parents presumably do not vaccinate because their limited financial resources make it difficult for them to follow the standard vaccination schedule. It is, however, somewhat puzzling why high-income parents do not vaccinate. One possibility is simply that higher-income parents have different information or beliefs. In this paper, however, I show that the hump shape can be explained even if everyone shares the same information and beliefs.[4] My explanation is tied to substitutes for vaccination. As parents' income increases, alternative options through which they can protect their children become available.[5] These can be either avoidance measures that reduce the risk of infection, or medical care that mitigates disease symptoms. If these alternative options are sufficiently effective—or so parents believe—, high-income parents will decide not to vaccinate their children.

Stated in this way, the point may seem obvious. Parents' vaccination decision problems are, however, more complicated because everyone's vaccination decisions are interdependent due to the positive externality involved in vaccination. To examine this mechanism more closely, I develop a simple static model of parents' vaccination decisions. To make my point as clearly as possible, I focus the analysis on medical care and assume that the choice of medical care is binary. Avoidance offers similar trade-offs but it may involve an additional complication known as fatalism in some cases (Kremer, 1996; Auld, 2003, 2006). I discuss the implication of avoidance in detail in a subsequent section.

The model has two key features. First, it assumes a population that consists of a continuum of agents ("parental units") who have different levels of income. Heterogeneity in income is necessary to examine the conditions under which the hump-shaped relationship can emerge within a population. The continuum assumption keeps the model tractable. Because the vaccination decisions of others change the probability of contracting a disease—and therefore agents' vaccination decisions—the population's vaccination rate is endogenously determined. Second, it focuses on the trade-off parents face in weighing the risk of side effects and disease. The side effects of vaccination may be mild such as pain and fever, or they may be more serious such as Guillain-Barre Syndrome and other autoimmune reactions. I assume the risks and consequences of side effects are common across all the agents. In contrast, medical care can lessen the duration of the disease or mitigate symptoms. This possibility

[3] Even though I include fixed effects and other control variables in the estimation, there may be a region-specific time-varying omitted variable (e.g., education) in the error term that is correlated with both income and vaccination rates. If so, it may not be income itself that causes vaccination rates to rise and fall. Although possible, it is somewhat difficult to imagine a variable (other than income) that systematically affects vaccination rates first positively and then negatively as it increases. Moreover, even if such a variable exists, vaccination rates continue to rise and fall along the same pattern outlined in this paper, as the variable increases.

[4] I do not exclude the possibility of misinformation and misbelief; people may share the same incorrect information and belief.

[5] Yang et al. (2016) mention such a possibility but do not provide a formal model to support their argument.

means that high-income parents may choose not to vaccinate their children if medical care is (believed to be) sufficiently effective.

Within this framework, I am able to show how the model can replicate the features observed in the data. The model clarifies the importance of distinguishing between two types of costs involved in vaccination: costs that are common to all the agents and opportunity costs that vary by agents. Common costs such as filling in forms, finding medical records, and transport to a clinic prevent low-income agents from vaccinating. In contrast, opportunity costs are key to understanding why high-income agents choose not to vaccinate. Both side effects and infection involve opportunity costs because an agent loses time in the occurrence of these events. As time is more valuable for higher-income agents, they choose an option with less expected time loss. This means that if medical care can effectively mitigate the duration or symptoms of the disease, high-income agents will choose not to vaccinate. The model shows under what conditions these individual vaccination decision rules drive equilibrium results.

The empirical literature on vaccination choice is extensive. While many studies find that low-income people are less likely to vaccinate (e.g., Wu et al., 2008; Klevens & Luman, 2001), some studies also find that high income is an obstacle for vaccination (e.g., Wei et al., 2009; Yang et al., 2016). Two studies make a step forward and compare the characteristics of parents who have an undervaccinated child with those who have an unvaccinated child (Smith et al., 2004), and the characteristics of parents who delay and refuse to vaccinate their child with those who only delay vaccination (Smith et al., 2011).[6] The findings of these studies suggest that the reason for not following the vaccination schedule may differ between low- and high-income groups. Thus, the literature is relatively successful in identifying the characteristics of these high-income and non-vaccinating parents. The literature, however, fails to clarify whether these parents are outliers or not. My contribution to this literature is to find the systematic hump-shaped relationship between income and vaccination rates.

Although the formal theoretical analysis of vaccination decisions under income heterogeneity I provide is novel, this paper draws on two branches of the economic epidemiology literature. One branch studies vaccination decisions in a population where agents differ in their cost of vaccination (e.g., Brito et al., 1991; Xu, 1999; Kureishi, 2009; Chen & Toxvaerd, 2014). These papers recognize that the cost of vaccination involves time loss in the event of side effects. It is, then, a natural extension to consider that the cost of infection also involves a similar time loss. By introducing income heterogeneity, this paper analyses agents' vaccination decisions when both the costs of vaccination and infection vary by agents through the difference in their value of time. The other branch studies the dynamics of infectious diseases in a population consisting of forward-looking agents (e.g., Francis, 1997; Goldman & Lightwood, 2002; Gersovitz & Hammer, 2003; Gersovitz, 2003; Barrett & Hoel, 2007; Toxvaerd, 2010a,b). One theme of this literature is to examine how prevention and treatment interact in disease management (e.g., Wiemer, 1987; Gersovitz & Hammer, 2004; Rowthorn & Toxvaerd, 2012). In my model, the substitution between vaccination and medical care,

[6] A child is categorised as "under-vaccinated" if he/she is not up-to-date on at least one vaccine but has received at least one dose of any of the six recommended vaccines, and is categorised as "unvaccinated" if he/has received no vaccinations (Smith et al., 2004).

interacted with the heterogeneity in income, generates a hump-shaped relationship between income and vaccination decisions.

Finally, the paper has an obvious connection to the Environmental Kuznets Curve (EKC) literature. Ever since Grossman & Krueger (1991, 1994) find evidence that environmental degradation initially rises and then falls as income increases, a number of empirical studies have tried to confirm its existence (e.g., Shafik & Bandyopadhyay, 1991; Cole et al., 1997; List & Gallet, 1999; Harbaugh et al., 2000). At the same time, four branches of theoretical explanations have been proposed: income effects with non-homothetic tastes (Lopez, 1994); threshold effects (John & Pecchenino, 1994; Stokey, 1998); increasing returns (Andreoni & Levinson, 2001); and technological progress coupled with neoclassical convergence (Brock & Taylor, 2010).[7] In this paper, I provide the first evidence that vaccination rates also initially rise and then fall as income rises, which could be labelled the "Vaccination Kuznets Curve (VKC)."[8] In contrast to the EKC, the VKC raises an alarm that economic development may bring falling vaccination rates and rising disease outbreaks. As such this paper is meant to be provocative rather than conclusive. I provide preliminary evidence and a suggestive explanation but certainly more work is warranted to identify the underlying mechanism at work.

The rest of the paper proceeds as follows. In section 2, I investigate the relationship between childhood vaccination rates and per capita income using country-, county-, and individual-level datasets. Section 3 develops a model that focuses on the role of medical care to explain the empirical regularity found in the previous section. Section 4 discusses alternative explanations for the empirical result with a special focus on avoidance measures. Section 5 concludes.

2 Empirical Analysis

This section examines the relationship between childhood vaccination rates and per capita income using three datasets with different levels of aggregation. In section 2.1, I use a country-level dataset collected by the WHO. I then proceed to county-level data from the U.S. in section 2.2. Finally, in section 2.3, an individual-level dataset from the U.S is analysed. As we proceed, the data coverage in terms of area and timespan becomes smaller, but more detailed information becomes available.

2.1 Country-level analysis

2.1.1 Data

The vaccination data is obtained from the WHO's Department of Immunization, Vaccines and Biologicals (IVB). The data is an estimate of the "infant" vaccination coverage rates,

[7]See the chapter 2 of Copeland & Taylor (2003) for a discussion of various effects.

[8]Troesken (2015) uses cross-country data to show that smallpox mortality and GDP per capita had a clear U-shape relationship at the beginning of the 20th century. This is indirect evidence that the Vaccination Kuznets Curve may have already existed a century ago.

and is presented as the percentage of a target population that has been vaccinated. For those vaccines given at birth like Bacille Calmette Guerin (BCG), the target population is the number of live births. For other infant vaccines such as diphtheria toxoid, tetanus toxoid and pertussis (DTP), the target population is children who survived their first birthday.[9]

The dataset covers nearly all the countries in the world, but the analysis uses only a subset of them. I exclude countries that are categorized as "low" and "lower middle" income countries by the World Bank. As the World Bank updates its country classification every year, I use the one in the year 2000, which is in the middle of the sample period (1980-2014). This procedure is to make sure that the sample does not include countries where the supply of vaccines is severely limited. In these countries, vaccination rates are presumably determined by limited supply and do not reflect choice behaviour. Data reliability is also a concern for these low-income countries. The final sample includes 70 countries during the period 1980-2014. Alternatively, dropping countries with the mean GDP per capita less than \$6,000, \$7,000, \$8,000 \$9,000 or \$10,000 gives a similar result.

Figure 1 depicts the relationship between vaccination rates and GDP per capita (hereafter per capita income) in the raw data for the following six vaccines: the first dose of BCG (BCG), the third dose of DTP (DTP3), the first dose of measles-containing vaccine (MCV), the third dose of Polio vaccine (Pol3), the third dose of Hepatitis B vaccine (HepB), and the third dose of Haemophilus influenzae type B vaccine (Hib3).[10] It shows a general pattern that vaccination rates are higher when per capita income is higher. The outliers in these graphs all originate from three specific countries. In BCG, DTP3, MCV and Pol3, the outliers with high per capita income (>\$50,000) and low vaccination rates (<60%) come from the United Arab Emirates. The outliers in BCG that have middle income (\$30,000-\$50,000) and low vaccination rates (<40%) are from Sweden and Ireland. Excluding these countries from the sample does not qualitatively change the results.

Although the positive relationship between vaccination rates and per capita income seems robust, it does not necessarily mean that these variables are directly related. Instead, it may merely mean that there is a confounding factor that is correlated with these variables. For example, countries with better institutions are likely to enjoy higher per capita income and higher vaccination rates at the same time. Then, the positive relationship between vaccination rates and per capita income may simply reflect the difference in institutional qualities across countries. To examine such possibilities, I conduct a regression analysis in

[9]Vaccination coverage rates are estimated based on two different sources of information: administrative data reported by each country and household surveys conducted by the WHO and United Nations Children's Fund (UNICEF). Because supplementary vaccinations received during occasional vaccination campaigns are not included, actual rates can be higher than the estimated numbers. It is important to note that the WHO and UNICEF estimate vaccination rates without any statistical analyses or modelling exercises. Rather, they subjectively determine vaccination coverage rates by examining the reliability of the reported and survey data. When the data is inconsistent or unreasonable, they make an attempt to identify the underling cause with help from local experts. When the data is missing in some years, they interpolate the data using data from other years in the same country (Burton et al., 2009). This procedure assures that there is no mechanical relationship between vaccination rates and per capita income.

[10]In addition to these "conventional vaccines", the dataset also includes seven other types of vaccines. Their data coverage is, however, quite limited both in terms of countries and years.

Figure 1: Vaccination rates and GDP per capita in raw data: Cross country data

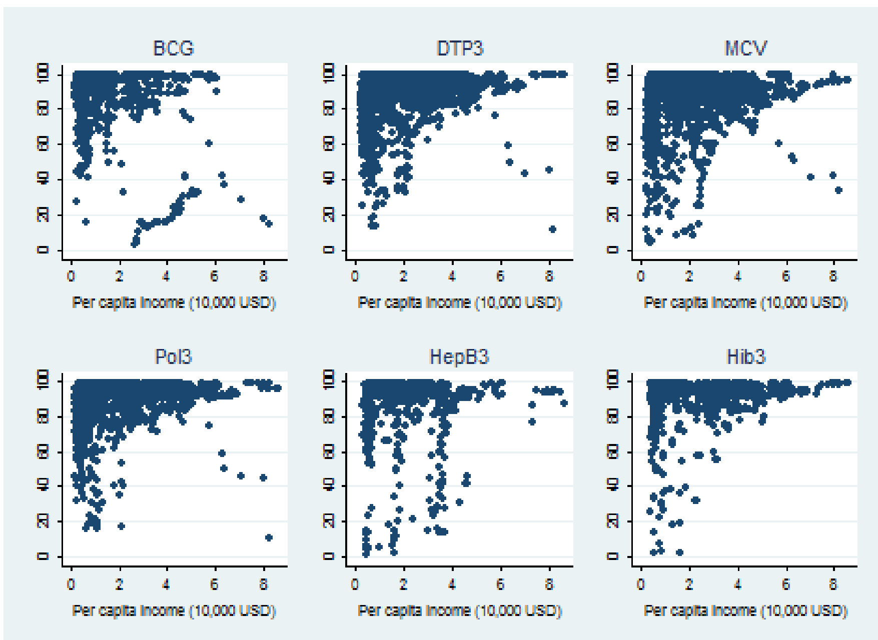

Source: World Health Organization
Note: BCG: The 1st dose of Bacille Calmette Guerin vaccine for tuberculosis, DTP3: The 3rd dose of diphtheria toxoid, tetanus toxoid and pertussis vaccine, MCV: The 1st dose of measles-containing vaccine, Pol3: The 3rd dose of Polio vaccine, HepB3: The 3rd dose of Hepatitis B vaccine, Hib3: The 3rd dose of Haemophilus influenzae type B vaccine.

the next section.

2.1.2 Econometric model

The econometric model is given by:

$$y_{it} = \alpha_i + \mu_t + \beta_1 G_{it} + \beta_2 G_{it}^2 + \beta_3 G_{it}^3 + \boldsymbol{X_{it}\gamma} + u_{it} \quad (1)$$

where y_{it} is the vaccination rate in country i at year t, G_{it} is per capita income, α_i and μ_t are respectively country- and year-fixed effects, $\boldsymbol{X_{it}}$ is a vector of controls, and u_{it} is the unobserved error term. The quadratic and cubic terms of income are included to capture any non-linear relationship between per capita income and vaccination rates. In addition to this cubic function regression, I also estimate a model with dummy variables corresponding to each 10,000 USD income range for per capita income. This dummy variable regression is more flexible than the cubic function regression, so the comparison between the two results will allow me to examine how well the cubic function approximates the underlying relationships.

The key in this equation is the inclusion of country- and year-fixed effects. Country-fixed effects control for any country-specific time-invariant factor such as institutions and culture. This is particularly important in this context, because institutional and cultural differences are substantial across countries. Year-fixed effects also play an important role because the dataset covers an extended period of time (i.e., 1980-2014).

Country- and year-fixed effects, however, do not control for factors that vary over time in a country-specific manner. If such factors are correlated with vaccination rates and per capita income, parameter estimates may still suffer from omitted variable bias. To mitigate this concern, I also include the following demographic characteristics of each country collected from the World Development Indicator as control variables in the equation: the total population, the shares of population aged 15-64 and over 65, the share of female population, and the share of rural population, and population density. These demographic characteristics are important determinants of a country's income. At the same time, demographic characteristics, and population density in particular, affects the risk of disease outbreaks, which in turn affects one's vaccination decisions. Therefore, they may become a source of omitted variable bias if left in the error term.

2.1.3 Summary statistics

Table 1 shows summary statistics of vaccination rates and per capita income. The mean vaccination rate is highest for Hib3, followed by Pol3 and DTP3. HepB3 and Hib3 have a relatively small number of observations because the dataset dose not include these vaccines in the 1980s. BCG also has a smaller number of observations because many developed countries do not use this vaccine any more.[11] Per capita income is obtained from the WDI.

[11]BCG is not effective in preventing primary infection. It is only effective in preventing TB patients from developing meningitis and disseminated TB. That is, BCG has little impact on TB transmission. Further, BCG may cause a false-positive reaction to a skin test for TB. See http://www.cdc.gov/tb/publications/factsheets/prevention/bcg.htm. For this reason, many developed countries such as the U.S. and U.K. do not include BCG in their routine vaccine schedule.

It is measured in market exchange rates and presented in $10,000 US in 2005.

Table 1: Summary statistics: Cross country data

Variable	N	Mean	SD	Min	Max
3rd dose of Pol	2034	89.46	12.53	11.00	99.00
3rd dose of DTP	2030	88.34	13.56	11.00	99.00
1st dose of BCG	1154	88.13	18.09	3.00	99.00
1st dose of MCV	1990	85.70	16.15	4.00	99.00
3rd dose of Hib	964	90.45	13.13	2.00	99.00
3rd dose of HepB	938	86.87	19.06	1.00	99.00
Per capita income	2178	2.028	1.555	0.18	8.61

Note: The dataset includes 70 countries during the period 1980-2014. Per capita income is presented in $10,000 in 2005. Pol: Polio vaccine, DTP: Diphtheria toxoid, tetanus toxoid and pertussis vaccine, BCG: Bacille Calmette Guerin vaccine, MCV: Measles-containing vaccine, Hib3: Haemophilus influenzae type B vaccine, HepB: Hepatitis B vaccine.

2.1.4 Result

The OLS estimates for each of the six vaccines are reported in Table 2.[12] Per income variables are separately not significant for most of the vaccines due to multicollinearity. They are, however, jointly significant at the conventional level. One exception is HepB3 with the p-value being 0.809. For each vaccine, the point estimates for the income variables exhibit both positive and negative signs, indicating a potential non-monotonic relationship between vaccination rates and per capita income. To see this, suppose that per capita income increases by $\$1,000$. For Pol3, other things being equal, this will increase the vaccination rate by 1.18 points if per capita income is $\$5,000$, but will decrease it by 0.07 points if per capita income is $\$30,000$. Similarly, for BCG, this will increase the vaccination rate by 0.81 points if per capita income is $\$5,000$, but will decrease it by 0.27 points if per capita income is $\$30,000$. As for the control variables, there is no clear pattern across vaccines as to which variables affect vaccination rates.

To understand the relationship between vaccination rates and per capita income, I calculate the predicted vaccination rates for an "average country" that takes the average value for the fixed effects and control variables. The predicted vaccination rate is given by (Grossman & Krueger, 1991):

$$\hat{y_{it}} = \bar{\alpha_i} + \bar{\mu_t} + \bar{\boldsymbol{X}}\hat{\boldsymbol{\gamma}} + \hat{\beta}_1 G_{it} + \hat{\beta}_2 G_{it}^2 + \hat{\beta}_3 G_{jt}^3 \tag{2}$$

where $\bar{\alpha_i}$ and $\bar{\mu_t}$ are respectively the average value of country- and year-fixed effects, $\bar{\boldsymbol{X}}$ is a vector of the average value of control variables, and $\hat{\beta}_i$ ($i = 1, 2, 3$) is the estimated parameter

[12] The full result including control variables is found in Appendix 6.1 and the result for the dummy variable regression is found in Appendix 6.2.

Table 2: Estimation results, Cross country

	Pol3	DTP3	BCG	MCV	Hib3	HepB3
Income	14.240	7.861	10.108	14.231	31.654^{b}	13.966
	(8.560)	(8.530)	(12.883)	(9.745)	(14.843)	(19.913)
Income squared	-3.165	-0.787	-1.979	-3.202	-8.209^{b}	-3.794
	(1.966)	(2.363)	(3.235)	(2.392)	(3.392)	(4.525)
Income cubed	0.149	-0.047	-0.037	0.186	0.531^{b}	0.283
	(0.153)	(0.200)	(0.236)	(0.182)	(0.229)	(0.309)
N	2034	2030	1154	1990	964	938
Countries	65	65	42	65	63	56
P(G=0)	0.003	0.053	0.000	0.031	0.042	0.809
R2	0.400	0.488	0.506	0.586	0.309	0.377
RMSE	8.232	8.027	7.705	9.074	9.169	12.004

Notes: Income is in $10,000. Standard errors are clustered at the country level. a, b, and c mean statistical significance at the 1 %, 5%, and 10% levels, respectively. All the specifications include country- and year-fixed effects, and demographic control variables. P(G=0) shows p-values for the joint test of income variables.

value. I also calculate the predicted vaccination rate from the dummy variable regressions. The relationships between the predicted vaccination rate and per capita income are presented in Figure 2. The solid lines show the predicted rates from the cubic function regression while circles show the predicted rates from the dummy variable regression. The size of the circles are proportional to the relative number of observations in each income range. The graphs show that the predicted vaccination rates follow a hump-shaped curve as per capita income increases for all the vaccines. The predicted rates from cubic and dummy regressions coincide very precisely, suggesting that the cubic function is a good approximation for the underlying relationship between vaccination rates and per capita income. Moreover, the size of these circles indicate that the hump shape is not simply driven by a small number of observations in the high-income range. The turning points for these hump shapes somewhat differ across vaccines: BCG, MCV, Pol3, HepB3 and Hib3 start decreasing at around $30,000 while DTP3 start declining at around $40,000. In percentage terms, the turning points are around 35-45% of the maximum income in the sample. This is a relevant range because the per capita income of many developed countries today falls into the income range of $30,000 - $40,000.

2.1.5 Short Discussion

The results indicate that childhood vaccination rates initially rise but then begin to fall as per capita income increases, conditional on fixed effects and control variables. It is tempting to interpret this result as causal: per capita income has positive effect on vaccination rates in the low-income range, but it has negative effect in the high-income range. There are, however, at least three reasons why such an interpretation may be incorrect: omitted variables, reverse causality, and measurement errors. First, even though I control for fixed effects and demographic characteristics in the estimation, it is still possible that the hump shape

Figure 2: Vaccination rates and GDP per capita conditional on fixed effects and controls: Cross country data

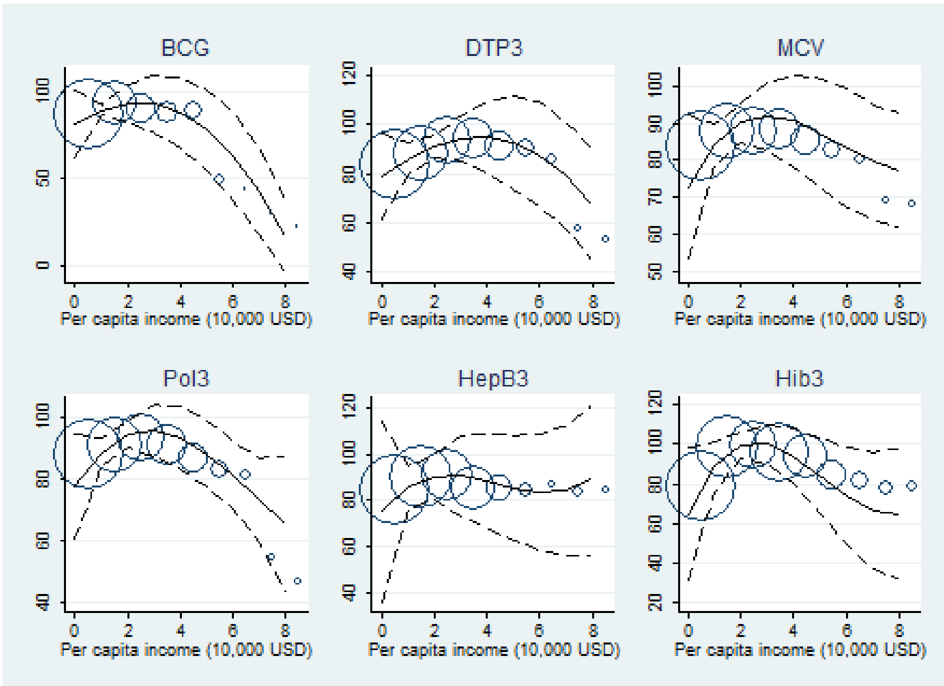

Note: Simulated vaccination rates are obtained by first regressing actual vaccination rates on income, fixed effects and controls, and then calculating predicted values fixing the value of these fixed effects and controls at their means. The solid line is the mean predicted rate while the dashed lines show 95% confidence interval from the regression with the cubic function of income. Each circle shows the predicted vaccination rate from the regression with dummy variables corresponding to each $10,000 income range. The size of each circle is proportional to the relative number of observations in each income range.

is driven by omitted factors that vary over time in a country-specific manner. Consider, for example, public health policies in each country. As countries become richer, the focus of public health policies typically shifts from infectious diseases to chronic diseases. This is because the incidence of infectious diseases falls due to nutritious food, sanitary environment, and intensive public health policies in addition to vaccination. This shift in public health policies may then generate an initial increase followed by a decrease in vaccination rates as per capita income rises.

Second, reverse causality from vaccination rates to per capita income may exist, either through the change in population or through the change in income. For example, a higher vaccination rate reduces children's mortality rate from the corresponding infectious disease. This will increase the population size, which in turn reduces per capita income for a given total income. Alternatively, a higher vaccination rate may increase the total income of a country because it also protects adults in the labour force. This will increase per capita income for a given population size. Therefore, in principle, the combination of these two channels can generate a hump-shaped relationship between per capita income and vaccination rates.

Finally, per capita income may be measured with errors. This is more likely in relatively low-income countries. If the true per capita income and the measurement error are correlated, the parameter estimates for per capita income will be biased toward zero; a so-called classical measurement error problem. Therefore, the relationship between per capita income and childhood vaccination rates may be at least partially masked by the measurement errors. Although this alone may not generate a hump shape, it is an obstacle in precisely examining the relationship between vaccination rates and per capita income.

To address these concerns, in the next section, I repeat the same analysis using the county-level data from the U.S. Although this dataset forces us to focus only on the U.S. in more recent years, it has the following advantages. First, public health policies are less diverse within the U.S. as compared to those across countries. Therefore, the omitted variable problem is unlikely to cause a serious bias. Second, the incidence of childhood diseases is extremely rare in the U.S. This makes the reverse causality channel less likely. Finally, per capita income within the U.S. is measured more accurately. This makes the measurement error problem unwarranted. With these advantages, the county-level dataset will give us a better idea about the relationship between childhood vaccination rates and per capita income.

2.2 County-level analysis

2.2.1 Data

The county-level vaccination data is originally from the National Immunization Survey (NIS), which is a random telephone survey to parents, followed by a survey sent to children's immunization providers. The survey began in 1994 and targets children between the ages of 19 and 35 months living in the U.S. Due to confidentiality reasons, the NIS public-use datasets do not include the information on the county of residence. Fortunately, using the NIS data, Smith & Singleton (2011) estimate vaccination rates in counties where the combined sample size from the NIS from at least one of the seven biennial periods during 1995 - 2008 is ≥ 35. This gives us a sample of 229 counties.

The dataset includes seven types of vaccines: the fourth dose of DTP (DTP4), the first dose of measles-mumps-rubella vaccine (MMR), the third dose of Pol (Pol3), the third dose of Hib (Hib3), the third dose of HepB (HepB3), the first dose of varicella vaccine (VRC), and the fourth dose of pneumococcal conjugate vaccine (PCV4). When compared to the country-level dataset in the previous section, several things are worth noting. First, this dataset does not include BCG because the U.S. does not require BCG in its standard vaccination schedule. Second, because the U.S. requires four doses of DTP vaccines instead of three, this dataset includes DTP4 as opposed to DTP3. Finally, while the country-level dataset includes MCV, the county-level dataset includes MMR, a specific type of MCV.

One potential caveat of this dataset is sample selection, because the 229 counties are not randomly selected from 3,141 counties in the U.S. This sample selection is an issue if these counties have some unobservable characteristics that are systematically correlated with vaccination rates. Fortunately, the 229 counties are selected purely based on the number of observations in the NIS, which is determined by the number of counties in each statistical area.[13] Then, even if the 229 counties have some unobservable characteristics that are correlated with vaccination rates, these characteristics must be time-invariant. Therefore, county-fixed effects will take care of any potential issue regarding the sample selection.

Figure 3 shows the relationship between vaccination rates and per capita personal income (hereafter per capita income) at the county level, for the seven vaccines.[14] It shows a general pattern that vaccination rates are higher when per capita income is higher. This pattern is particularly clear for the income range $\$20,000 - \$60,000$, but it is less clear above this range. The reason is simply that there are only a small number of observations above $\$60,000$. Most of these are sampled from from New York county, in the state of New York. With the exception of this county, the graph shows a clear pattern: vaccination rates rise as per capita income increases.

Again, the relationship observed in the raw data may be confounded by other factors such as institutional differences across counties or gradual changes in the recognition of vaccine importance over time. To examine such possibilities, I proceed to a regression analysis.

2.2.2 Econometric model

The econometric model is given by:

$$y_{it} = \alpha_i + \mu_t + \beta_1 G_{it} + \beta_2 G_{it}^2 + \beta_3 G_{it}^3 + \boldsymbol{X_{it}\gamma} + u_{it} \tag{3}$$

where y_{it} is the vaccination rate in county i at year t, G_{it} is per capita income, α_i and μ_t are respectively county- and year-fixed effects, $\boldsymbol{X_{it}}$ is a vector of controls, and u_{it} is the unobserved error term. County- and year-fixed effects are again the key in this equation. County-fixed effects control for the difference in institutions and culture across counties, while year-fixed effects control for the gradual change in the public interest in vaccination over time. To control for other factors that vary over time in a county-specific fashion, I also include in the estimation the total population, the shares of people between 14-65 and above 65, the share of male population, the shares of white, black and native Indian population, and population density. These variables are collected from the Census.

Figure 3: Vaccination rates and per capita income in raw data: US county data

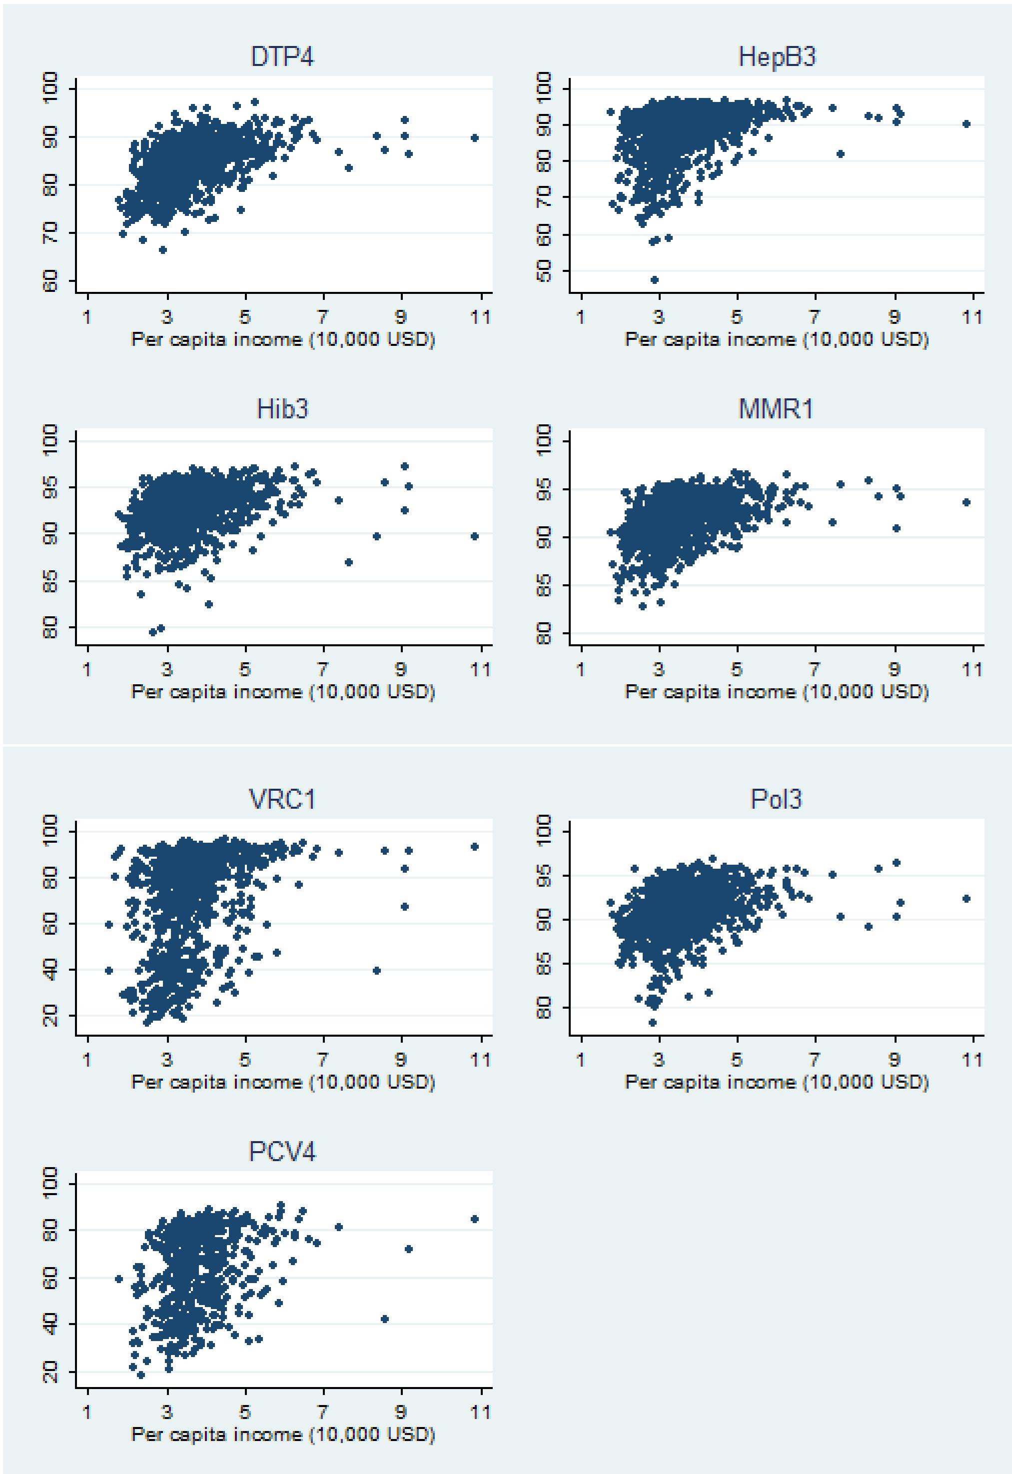

Source: Smith & Singleton (2011)

Note: The dataset includes 229 counties during the period 1995-2008. DTP4: The 4th dose of diphtheria toxoid, tetanus toxoid and pertussis vaccine, MMR1: The 1st dose of measles-mumps-rubella vaccine, Pol3: The 3rd dose of Polio vaccine, HepB3: The 3rd dose of Hepatitis B vaccine, Hib3: The 3rd dose of Haemophilus influenzae type B vaccine, VRC1: The 1st dose of Varicella vaccine, PCV4: The 4th dose of Pneumococcal Conjugate vaccine.

Table 3: Summary statistics: US country-level data

Variable	N	Mean	SD	Min	Max
3rd dose of Pol	1287	90.94	2.64	78.10	96.80
4th dose of DTP	1287	84.18	4.40	66.20	97.00
1st dose of MMR	1287	91.80	2.13	82.80	96.60
3rd dose of Hib	1287	92.81	2.20	79.20	97.30
3rd dose of HepB	1287	88.19	6.62	47.10	96.60
1st dose of VRC	1099	72.00	20.75	16.10	96.50
4th dose of PCV	556	61.31	17.11	17.70	91.00
Per capita income	2946	3.562	0.874	1.472	10.84

Notes: The dataset includes 229 counties during the period 1995-2008. Per capita income is presented in $10,000 in 2005. Pol: Polio vaccine, DTP: Diphtheria toxoid, tetanus toxoid and pertussis vaccine, MCV: Measles-containing vaccine, Hib: Haemophilus influenzae type B vaccine, HepB3: Hepatitis B vaccine, VRC: varicella (chickenpox) vaccine, PCV: pneumococcal conjugate vaccine.

2.2.3 Summary statistics

Table 3 shows summary statistics of the county-level dataset. The average vaccination rates are higher for most of the vaccines in the county-level dataset, when compared to the country-level dataset in the previous section. The two vaccines that are not included in the country-level dataset, VRC and PCV4, have a relatively low vaccination rate and a high variance. While per capita income in the country-level dataset ranges between $1,800 - $86,000, per capita income in the county-level dataset ranges from $15,000 to $110,000. Therefore, as compared to the country-level dataset, the county-level dataset include less observations in the low income range while it includes more observations in the high income range. Another thing worth noting is that per capita income in this dataset is much more concentrated. This can be seen by comparing the coefficient of variation $0.874/3.562 = 0.245$ in this dataset in contrast to that of $1.555/2.028 = 0.766$ in the country-level dataset. This is partly because the dataset is a small sub-sample of the entire counties.

[13]The NIS ensures that a similar number of samples are obtained in each of the 78 Immunization Action Plan areas (50 states, Washington DC, and 27 other large urban areas). Each area, however, has a different number of counties, ranging from only 3 counties in Delaware to 254 counties in Texas. This suggests that the counties in Delaware are included in the NIS sample every year while the majority of the counties in Texas are not included in the NIS sample for a given year. As a result, the 229 counties (out of 3,141 counties) in our dataset do not represent the entire number of counties in the U.S. Rather, they over-represent the Immunization Action Plan by favouring areas with a smaller number of counties.

[14]The personal per capita income is obtained from Bereau of Economic Analysis. It is constructed from the "personal income" rather than the national income. Note, however, that personal income and national income take similar values at the country level. According to http://thismatter.com/economics/national-accounts.htm, the main difference between personal income and national income is that personal income includes transfer payments, such as private pension payments, retirement benefits, unemployment insurance benefits, veteran benefits, disability payments, welfare, and farmer subsidies.

2.2.4 Results

The estimation results are found in Table 4.[15] Per capita income variables are jointly significant for all the vaccines at the conventional level. Two exceptions are Pol3 and Hib3. While Hib3 is significant at the 0.15 level (p-value: 0.114), Pol3 is not significant at all (p-value: 0.916). Again for all the vaccines, the point estimates for income variables exhibit both positive and negative signs, indicating a potential non-monotonic relationship between vaccination rates and per capita income. To see this, suppose per capita income increases by $1,000. For MMR1, other things being equal, this will lead to an increase in the vaccination rate by 0.17 points when per capita income is $20,000, but will lead to a decrease by 0.1 points when per capita income is $50,000. Similarly, for HepB3, this will lead to an increase in the vaccination rate by 0.39 points when per capita income is $20,000, but will lead to a decrease by 0.24 points when per capita income is $50,000.

Table 4: Estimation results, US county

	Pol3	DTP4	MMR1	Hib3	HepB3	VRC1	PCV4
Income	-0.538	-0.018	5.000^{a}	-5.056^{c}	11.432^{a}	21.726^{a}	46.622^{a}
	(2.386)	(3.610)	(1.851)	(2.742)	(3.688)	(7.706)	(16.645)
Income squared	0.056	0.090	-0.981^{a}	0.790^{c}	-2.203^{a}	-3.987^{a}	-7.962^{a}
	(0.399)	(0.530)	(0.310)	(0.470)	(0.614)	(1.252)	(2.747)
Income cubed	-0.001	-0.000	0.050^{a}	-0.036	0.108^{a}	0.187^{a}	0.382^{a}
	(0.021)	(0.025)	(0.016)	(0.025)	(0.031)	(0.061)	(0.132)
N	1287	1287	1287	1287	1287	1099	556
Counties	229	229	229	229	229	229	229
P(G=0)	0.916	0.046	0.004	0.114	0.000	0.000	0.040
R2	0.470	0.353	0.236	0.328	0.847	0.958	0.942
RMSE	1.536	2.251	1.367	1.380	2.436	3.966	3.633

Notes: Income is in $10,000. Standard errors are clustered at the county level. a, b, and c mean statistical significance at the 1 %, 5%, and 10% levels, respectively. All the specifications include county- and year-fixed effects, and demographic control variables. P(G=0) shows p-values for the joint test of income variables.

Figure 4 shows the predicted vaccination rates for an average county. The predicted vaccination rates show an inverted U-shape for HepB3, MMR1, VRC1 and PCV4. The turning points for these vaccines are around $40,000. In percentage terms, the turning points are around 35% of the maximum income in the sample, which is similar to those found in the country-level analysis. It is notable that there are many observations around the income of $40,000 as shown in Figure 3, which means the decrease in vaccination rate is not entirely driven by a small number of counties such as New York county.[16] Alternatively, DTP4 monotonically increases with income and Hib3 show a U-shape curve while Pol3 shows a flat curve. Therefore, the results are not as consistent across vaccines when compared to those in the country-level analysis in the previous section. Even so, it is assuring that four out of

[15]The full result including control variables is found in Appendix 6.3.

[16]Dropping New York county does not qualitatively change the result.

seven vaccines show a hump-shaped relationship with per capita income, the same pattern that was found at the country level.

2.2.5 Short Discussion

In this section, there is again a hump-shaped relationship between childhood vaccination rates and per capita income for more than half of the vaccines analysed. As noted earlier, the county-level analysis is less likely to suffer from omitted variable bias, reverse causality, and measurement errors, when compared to the country-level analysis in the previous section. Therefore, the results in this section are stronger evidence that per capita income does *affect* childhood vaccination rates in a non-monotonic fashion.

It is, however, still possible that the results in this section are driven by omitted variables. Recall that a potential source of omitted variable bias in the country-level analysis is public health policies. This is less of a concern in this section, because public health policies are presumably much less diverse with the U.S. After controlling for county- and year-fixed effects, there is much less scope for public health policies to cause omitted variable bias. This does not, however, guarantee that the bias does not exist. To the extent that such policies vary over time in a county-specific manner, they may still cause bias in the estimation.

Another set of potential omitted variable is demographic characteristics of each county. Both in the country- and county-level analysis, I control for demographic characteristics because they are potentially correlated with vaccination rates and per capita income. It is, however, still possible that some important population characteristics are left in the error term. To see this, consider parents who make vaccination decisions for their children. It is natural to assume that their vaccination decisions are affected by various factors, such as education, age, race, and the number of children. It is then clear that estimations at the country- and county-level analysis consider only a part of these factors. It is possible that these demographic factors that are not considered cause some bias.

To further address these concerns for the omitted variable bias, I repeat the same analysis again using an individual-level dataset from the U.S. in the next section. Because the individual-level dataset includes a large number of observations in each state, it is possible to add state-year-fixed effects in the estimation.[17] This is quite appealing because any public health policies that are specific to each state in each year will be controlled for. Also, the individual-level dataset includes detailed family characteristics. Controlling for these variables in the estimation further mitigates the concern for omitted variable bias.

2.3 Individual-level analysis

2.3.1 Data

The individual vaccination data again comes from the NIS. The data is available from the year 1994 but the analysis in this section uses data from the year 2005 when a key detailed

[17]It is actually possible to include state-year-fixed effects in the county-level analysis, which gives qualitatively similar results. Due to the smaller sample size, however, the estimation results are statistically weaker.

Figure 4: Vaccination rates and per capita income conditional on fixed effects and controls: US county data

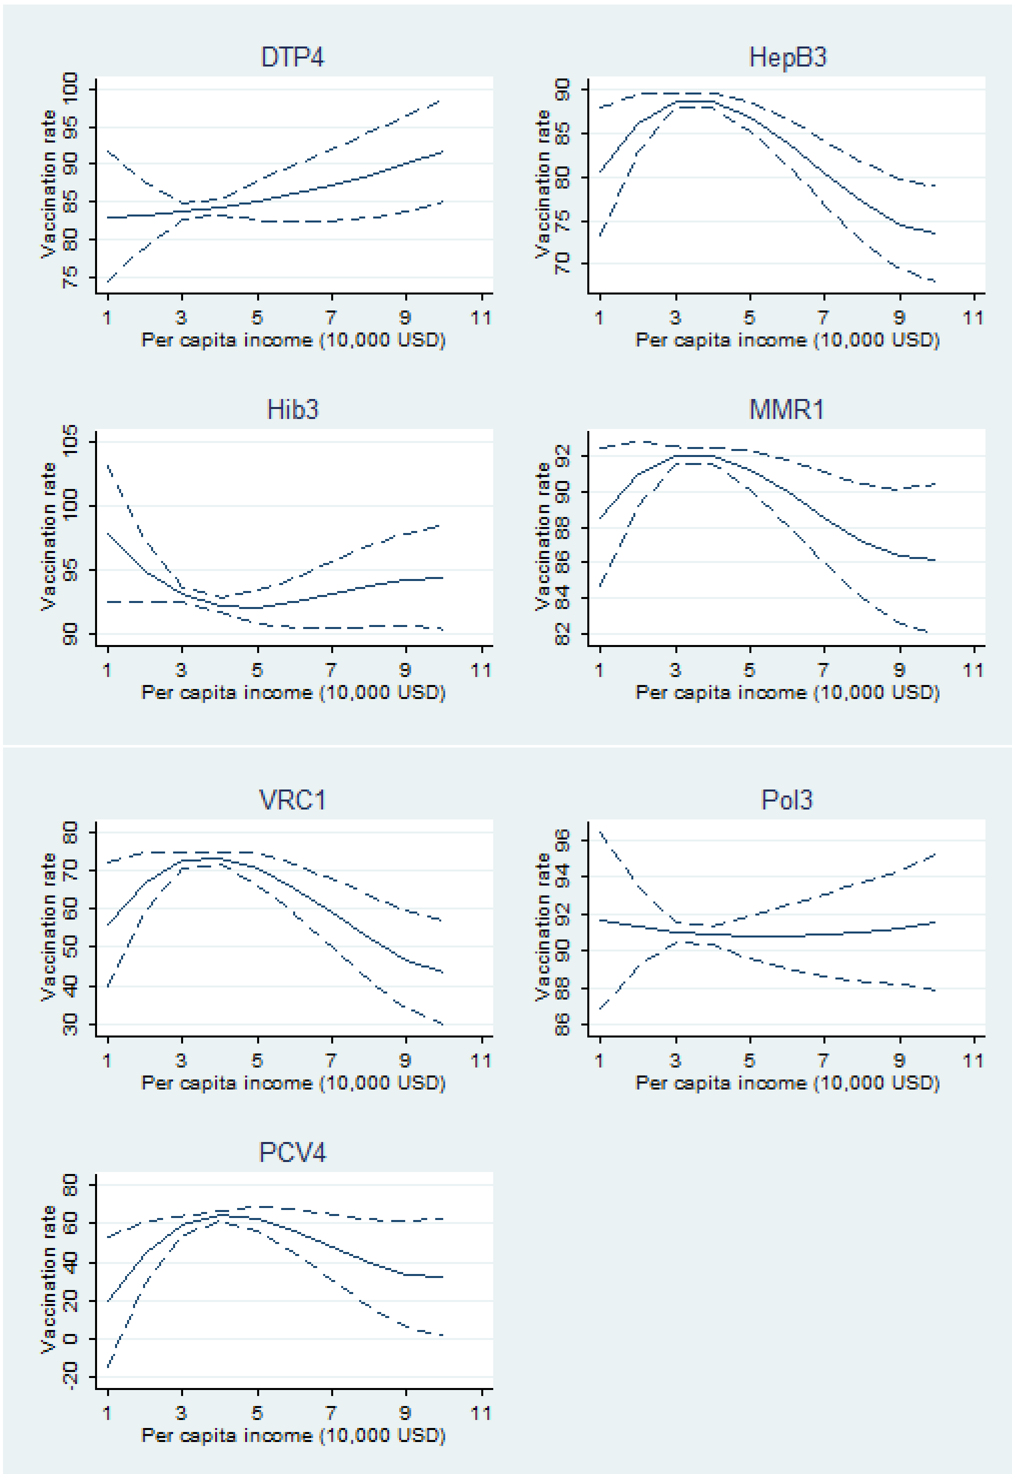

Note: Simulated vaccination rates are obtained by first regressing actual vaccination rates on income, fixed effects and controls, and then calculating predicted values fixing the value of these fixed effects and controls at their means. The solid line is the mean predicted rate while the dashed lines show 95% confidence interval.

income variable became available. The NIS includes the vaccination status of children for nine different vaccines: DTP, Pol, MCV, Hib, HepB, VRC, PCV, ROT, and FLU.[18] FLU is, however, excluded from the analysis because the recommended number of doses for young children changes from season to season and the up-to-date variables for flu vaccines in the dataset do not accurately reflect the latest recommendation.[19] As a result, the dataset in this section includes the same seven vaccines as the county-level dataset in the previous section. Note however that the individual-level dataset analyses MCV while the county-level dataset in the previous section analyses MMR. In addition to the seven vaccines, the first does of ROT is also analysed in this section, which is not included in the county-level dataset due to confidentiality reasons.

To conduct the same analysis as in the previous sections, I need the information about vaccination rates and per capita income. As the NIS provides individual-level data, the dependent variable in this section is a binary variable that assigns the value of 1 if a given child is up-to-date with the standard schedule for a vaccine at the age of 19-35 months and 0 if not. Turning to the independent variable, the income variable provided in the NIS is not per capita income, but family income. Moreover, it is provided as a range rather than an exact level, with the top income category censored. Therefore, I use the following three steps to construct per capita income. First, I assign each family the lower boundary of the income bin as their family income.[20] Then, I divide this family income by the number of family members to obtain per capita income. Finally, this nominal per capita income is converted into real per capita income in $2005 US using the consumer price index.

Figure 5 shows the relationship between the fraction of children who are up-to-date with a vaccine ("vaccination rate") and per capita income in each income range. The graph suggests that the fraction of children being up-to-date increases with per capita income in the low income range ($0 - $20,000) for all the vaccines. In the high income range ($20,000-$35,000), however, this tendency is less clear. In fact, for DTP4, HiB3 and PCV4, the fraction decreases as per capita income increases. Alternatively, for VRC1, the fraction keeps increasing with per capita income.

These patterns again do not necessarily mean that income and vaccination rates are directly related. Rather, there may be a third variable that is correlated with both vaccination rates and per capita income. For example, the residence of certain states earn higher income than the residence in other states. If these high-income states are the ones that enforce more stringent policies against childhood vaccination exemption, this will create a spurious positive correlation between income and vaccination rate. To eliminate such possibility, I conduct a regression analysis.

[18]ROT: rotavirus vaccine, FLU: seasonal influenza vaccine.

[19]See "A Codebook for the 2014 Public-Use Data File" in the National Immunization Survey.

[20]Family income is given as 12 income categories: $0-$7,500; $7,500-$10,000; $10,000-$17,500; $17,500-$20,000; $20,000-$25,000; $25,000-$30,000; $30,000-$35,000; $35,000-$40,000; $40,000-$50,000; $50,000-$60,000; $60,000-$75,000 and $75,000+. The lower boundary was chosen because the highest income group does not have the upper boundary or midpoint.

Figure 5: Probability of being up-to-date with a vaccine and per capita income in raw data: US individual data

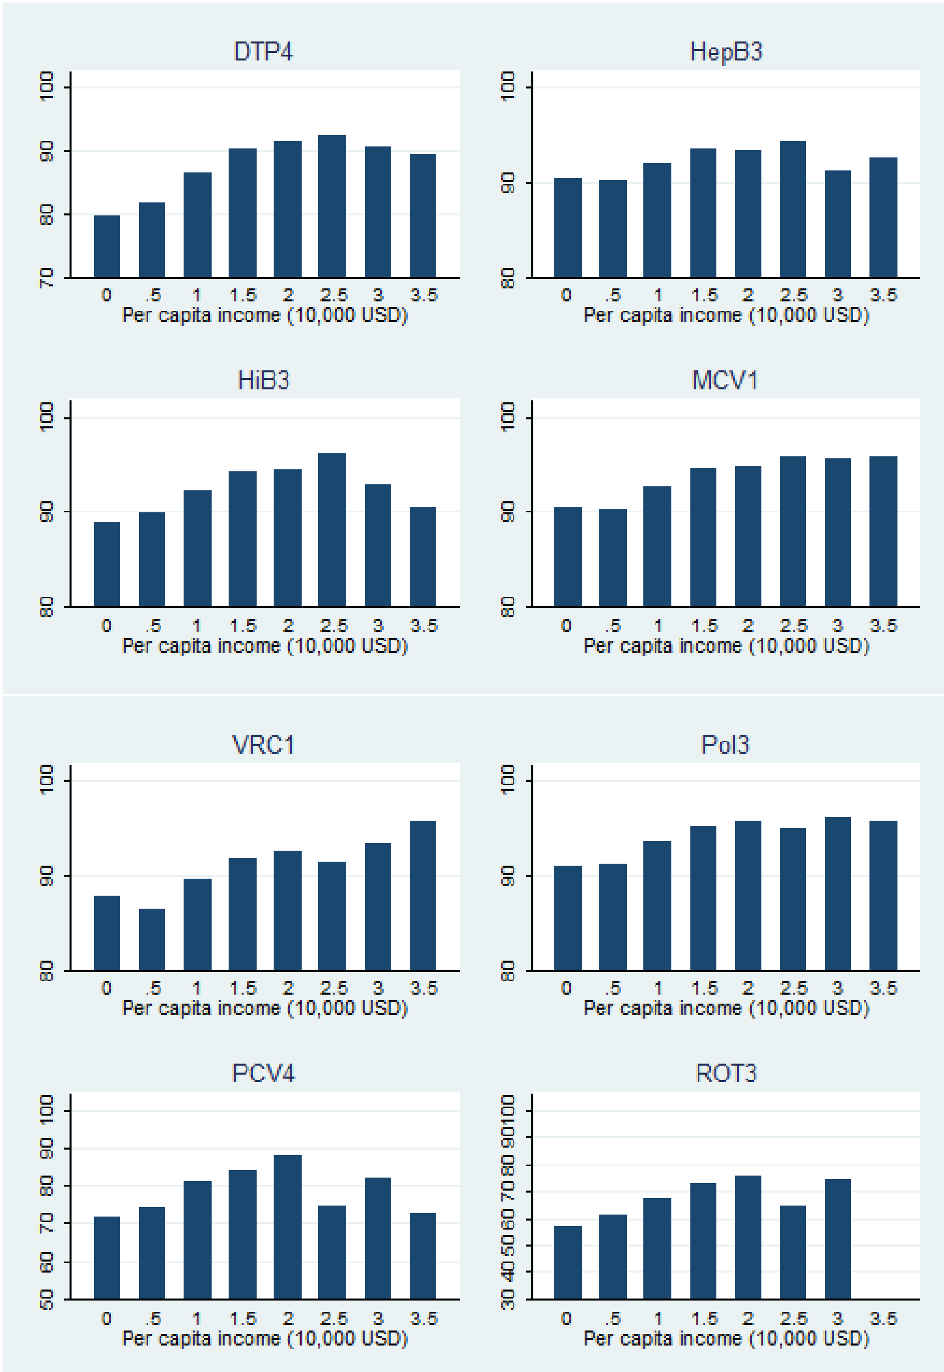

Source: CDC, NCRID and NCHS (2006-2015), 2005-2014 National Immunization Survey.
Note: DTP4: The 4th dose of diphtheria toxoid, tetanus toxoid and pertussis vaccine, MCV1: The 1st dose of measles-containing vaccine, Pol3: The 3rd dose of Polio vaccine, Hep3: The 3rd dose of Hepatitis B vaccine, HiB4: The 4th dose of Haemophilus influenzae type B vaccine, VRC1: The 1st dose of Varicella vaccine, PCV4: The 4th dose of Pneumococcal Conjugate vaccine, ROT3: The third dose of Rotavirus vaccine. Per capita income is calculated as family income divided by the number of family members.

2.3.2 Econometric model

I use the linear probability model given by[21]:

$$y_{ijt} = \alpha_{jt} + \beta_1 G_{ijt} + \beta_2 G_{ijt}^2 + \beta_3 G_{ijt}^3 + \boldsymbol{X_{ijt}}\boldsymbol{\gamma} + u_{ijt} \tag{4}$$

where y_{ijt} is a binary variable indicating whether a child i in state j is up-to-date with a vaccine at year t, G_{ijt} is per capita income, α_{jt} is the state-year-fixed effect, $\boldsymbol{X_{ijt}}$ is a vector of controls, and u_{ijt} is the unobserved error term. As the data is a repeated cross-section data, no individual-level fixed effects are included.

The inclusion of state-year-fixed effects is key in this estimation because they will control for any state-level public health policy in each year that can affect vaccination rates. Given that many public health policies are set at the state level, the inclusion of state-year-fixed effects virtually eliminate the concern for the omitted variable bias arising from public health policies. In addition, they also control for any demographic characteristics. For example, large cities typically have higher per capita income and higher population density at the same time. Then, population density is a source of omitted variable bias even after controlling for individual characteristics. State-year-fixed effects are extremely powerful in this sense, because they can control for any state-level demographic factor in each year.

In addition to state-year-fixed effects, a number of individual-level characteristics are controlled in the estimation. These characteristics may affect vaccination decisions and per capita income in various ways. For example, consider the age of the mother. It is natural that the age is positively related to income. At the same time, a higher age may mean more experience as a mother and more knowledge about vaccination, which will affect her vaccination decisions for her children. A similar logic can be applied to other characteristics. Therefore, I control for household characteristics (number of household members, number of children less than 18 years old, language with which interview was conducted, state of residence of child at birth versus current state, relationship of respondent to child), mother's characteristics (education, age, marital status) and child's characteristics (age, gender, first born status, race/ethnicity, Hispanic origin).

2.3.3 Summary statistics

Table 5 shows summary statistics of the individual data. The dataset covers the period 2005-2014 and includes more than 160,000 observations. The mean rates vary across vaccines, ranging from >90 percent for MCV and Pol to 66 percent for ROT. This seems to partially reflect the fact that the rota virus vaccine was introduced in 2006 and the up-to-date variable for the vaccine became available in 2009 in the public-use data file. ROT vaccine has a very high variation while MCV and Polio vaccine have a relatively small variation. The lowest value of per capita income is \$0, which reflects the fact that I assign 0 income to families in the lowest income bin. The highest value of per capita income is \$37,500, which reflects the

[21] I prefer the linear probability model because Probit and Logit models may suffer from the incidental parameter problem. Probit and Logit models that include only state- and year-fixed effects (but not state-year-fixed effects) give very similar results.

fact that the family income is top-coded at $75,000 and each family has at least one parent and one child. This is a limitation of this dataset because it does not precisely identify the observations around and above the income level $40,000 at which, according to the county-level analysis, vaccination rates begin to decline.

Table 5: Summary statistics: US individual-level data

Variable	N	Mean	SD	Min	Max
3rd dose of Pol	161002	93.30	24.99	0	100
4th dose of DTP	161002	85.48	35.22	0	100
1st dose of MCV	161002	92.58	26.20	0	100
3rd dose of Hib	161002	91.93	27.22	0	100
3rd dose of HepB	161002	92.01	27.11	0	100
1st dose of VRC	161002	89.61	30.50	0	100
4th dose of PCV	161002	78.68	40.95	0	100
3rd dose of ROT	92501	65.67	47.47	0	100
Per capita income	161002	10.37	6.60	0	37.5

Notes: The dataset includes the period 2005-2014. Per capita income is presented in $10,000 in 2005. Pol: Polio vaccine, DTP: diphtheria toxoid, tetanus toxoid and pertussis vaccine, MCV: Measles-containing vaccine, Hib: Haemophilus influenzae type B vaccine, HepB3: Hepatitis B vaccine, VRC: Varicella (chickenpox) vaccine, PCV: Pneumococcal Conjugate Vaccine, ROT: rotavirus vaccine.

2.3.4 Results

The estimation results are shown in Table 6. The income variables are jointly significant at the conventional level for all the vaccines. Again, the point estimates for the income variables exhibit both positive and negative signs, indicating a potential non-monotonic relationship between the probability of beingup-to-date and per capita income. For example, suppose again per capita income increases by $1,000. For DTP4, other things being equal, this will increase the probability by 0.24 points if per capita income is $15,000, but will decrease the probability by 0.28 points if per capita income is $30,000. Similarly, for HiB3, this will increase the probability by 0.10 points if per capita income is $15,000, and will decrease the probability by 0.27 points if per capita income is $30,000. In terms of the control variables, they indicate that the probability of a child being up-to-date is higher if the child is older, if the number of children in the family is smaller, if the interview was conducted in Spanish,[22] if the mother has a higher education, and if the family lives in the same place as the child was born.[23] It is particularly notable that the mother's education has a monotonic relationship

[22]Fry (2011) also reports that children whose interview was conducted in Spanish are more likely to be up-to-date with pertussis-containing vaccines.

[23]The full result including control variables is found in Appendix 6.4.

with the probability for all the vaccines; the higher the mother's education, the higher the probability of the child being up-to-date with every single vaccine.

Table 6: Estimation results, US individual

	Pol3	DTP4	MCV1	HiB3	HepB3	VRC1	PCV4	ROT3
Income	0.141	2.553	0.333	2.997^b	0.271	-1.044	6.271^a	13.833^a
	(1.363)	(1.877)	(1.356)	(1.506)	(1.393)	(1.462)	(2.103)	(3.169)
Income squared	0.316	0.819	0.476	-0.477	0.922	1.673	-0.422	-5.849^b
	(1.089)	(1.478)	(1.084)	(1.229)	(1.141)	(1.168)	(1.696)	(2.717)
Income cubed	-0.037	-0.382	-0.114	-0.108	-0.280	-0.340	-0.248	1.029
	(0.256)	(0.341)	(0.255)	(0.297)	(0.278)	(0.275)	(0.404)	(0.671)
N	161002	161002	161002	161002	161002	161002	161002	92501
P(G=0)	0.044	0.000	0.014	0.000	0.005	0.000	0.000	0.000
R2	0.024	0.055	0.024	0.045	0.020	0.025	0.096	0.107

Notes: Income is in $10,000. a, b, and c mean statistical significance at the 1 %, 5%, and 10% levels, respectively. All the specifications include state-year-fixed effects and other control variables. P(G=0) shows p-values for the joint test of income variables.

The relationships between the predicted probability of being up-to-date and per capita income are presented in Figure 6. The graph shows that the mean predicted probability of being up-to-date monotonically increases with income for Pol3 and ROT3 for the entire income range. Alternatively, for other six vaccines, it shows that the probability of being up-to-date initially rises but then falls as per capita income increases. The turning points are around $20,000-$30,000. As the income data is top-coded in this dataset, the turning points are not directly comparable to those in the country- and county-level analysis. The upper bound of the 95% confidence interval, however, actually rises as income increases, making it hard to draw a decisive conclusion about the behaviour of the probability in the high income range. This comes partly from the fact that the income is top-coded in this dataset so that I do not have many observations in the high income range. Yet, the fact that a hump shape is observed in six out of eight vaccines is suggestive that the probability decreases in the high income range.

2.4 Summary of the empirical findings

In section 1, I used three datasets to examine the relationship between per capita income and childhood vaccination rates. I started with the country-level dataset, which includes a large number of countries over an extended period of time. I found that all the vaccines revealed a hump-shaped relationship with per capita income once country- and year-fixed effects are controlled. The result indicates the possibility that income has positive effect on vaccination rates in the low-income range, but has negative effect in the high-income range. The interpretation of the results, however, requires caution due to the potential omitted variable bias, reverse causality, and measurement errors.

To deal with these concerns, I then proceeded to the county-level data from the U.S. Although this dataset limits the analysis to the U.S. in recent years, it greatly reduces the

Figure 6: Probability of being up-to-date with a vaccine and per capita income conditional of fixed effects and controls: US individual data

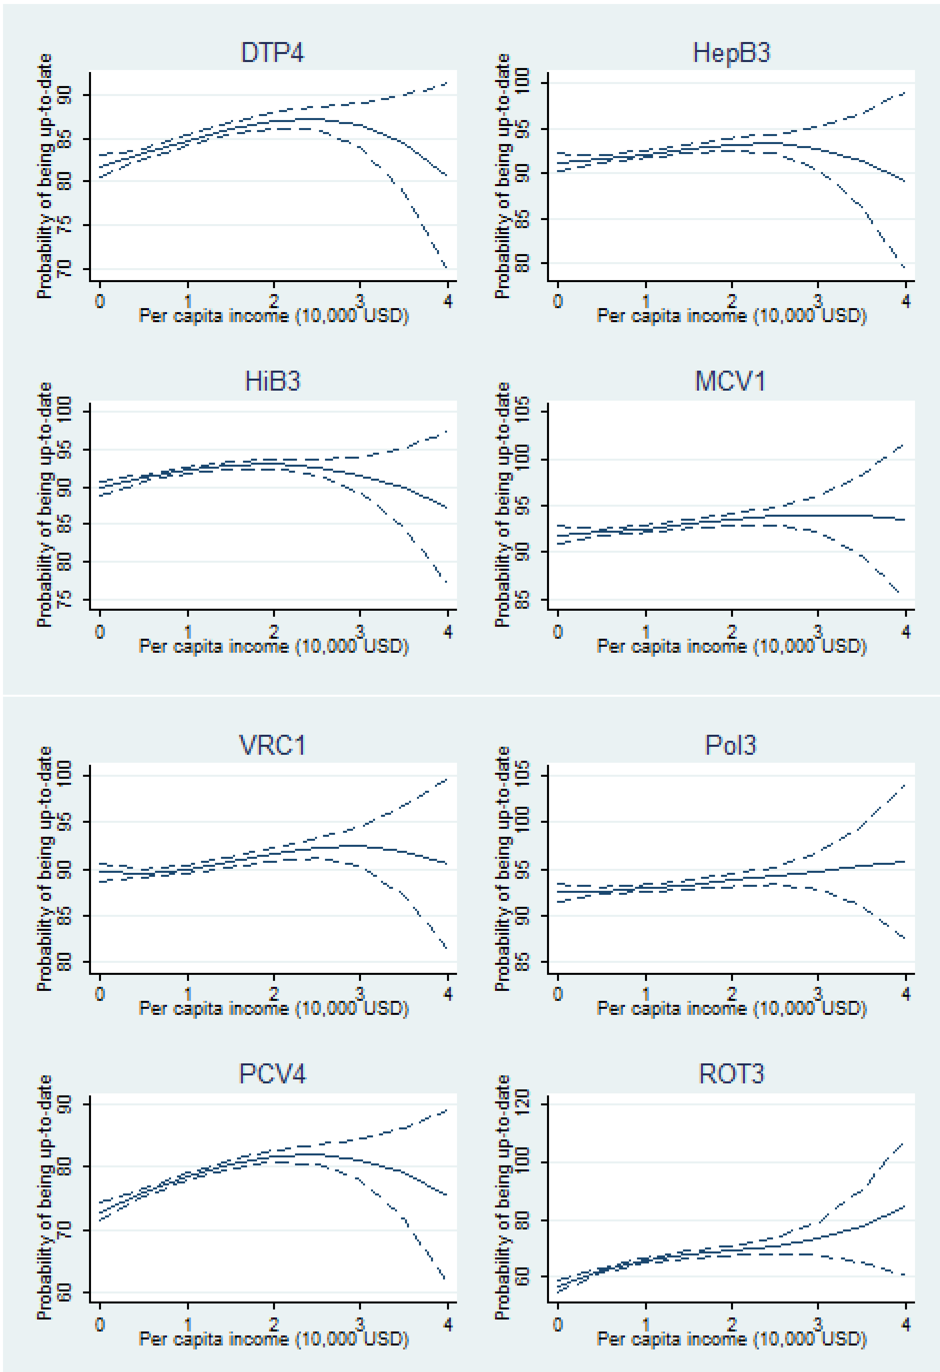

Note: Simulated vaccination rates are obtained by first regressing actual vaccination rates on income, fixed effects and controls, and then calculating predicted values fixing the value of these fixed effects and controls at their means. The solid line is the mean predicted rate while the dashed lines show 95% confidence interval.

concern for omitted variable bias, and virtually eliminates the concerns for reverse causality and measurement errors. After controlling for county- and year-fixed effects, I find that more than half of the vaccines showed a hump-shaped relationship with per capita income. The consistent results between country- and county-level analyses suggested that per capita income affects vaccination rates in a non-monotonic fashion.

To further reduce the concern for the omitted variable bias, I also analysed the individual-level data in the U.S. This dataset allows for the inclusion of state-year-fixed effects, which virtually eliminated the concern for omitted variable bias. The dataset has, however, one disadvantage in that the income is top-coded so that not many observations are available in the range where vaccination rates begin to fall. In this dataset, some vaccines exhibit a hump-shaped relationship with per capita income even in the raw data. After controlling for state-year-fixed effects, most of the vaccines revealed a similar hump-shaped relationship. This is strong evidence that supports the causal link from income to vaccination rates.

In sum, even though each dataset has its own limitations, the results are fairly consistent across vaccines and datasets. These results suggest that per capita income has positive effect on vaccination rates in the low income range, but it has negative effect on vaccination rates in the high income range. Moreover, the results indicate that the mechanism stems from individual-level behaviours: both low- and high-income parents are less likely to follow the standard vaccination schedule.

3 The Model

It appears that both low- and high-income parents are less likely to follow the standard vaccination schedule for children. Low-income parents presumably do not vaccinate because their limited financial resources make it difficult for them to follow the standard vaccination schedule. It is, however, somewhat puzzling that high-income parents do not vaccinate. One potential explanation is that there are substitutes for vaccination. The substitute could be either avoidance that reduces the risk of infection, or medical technology that effectively mitigates disease symptoms. If this substitute is effective enough—or so parents believe—then high-income parents may decide not to vaccinate or, at least, delay vaccinating their children. As everyone's vaccination choices are interdependent, however, it is not necessarily clear whether such a mechanism can work. To examine this possibility, in this section, I build a model of the vaccination decision. To make my point as clear as possible, I focus only on the variable of medical care rather than avoidance, and I assume the choice of medical care is binary.

3.1 The vaccination decision

Consider the problem of an agent who faces the risk of being infected with a non-fatal virus. The agent has an exogenous income y, which is the product of the agent's labour productivity y and the fixed work time 1. The agent decides whether to vaccinate or not, and spends the rest of the income for consumption. Vaccination costs c_v and gives the agent perfect immunity, but it may also involve side effects with a probability q. In the event of side effects, the agent loses a fraction α of work time, which reduces the agent's income to $(1-\alpha)y$. Then, with vaccination, the agent's expected utility $u_v(y)$ is given by:

$$u_v(y) = q * u((1-\alpha)y - c_v) + (1-q) * u(y - c_v) \tag{5}$$

Without vaccination, the agent will contract the disease with a probability p. This probability is perceived as given by each agent, even though it is endogenously determined by the vaccination rate in the population. In the event of an infection, the agent loses a fraction γ_0 of work time, which reduces the agent's income to $(1-\gamma_0)y$. If the agent uses medical care with the cost c_t, however, the loss in work time is reduced such that $\gamma_1 < \gamma_0$. Therefore, without vaccination, the agent's expected utility $u_{nv}(y)$ is given by:

$$u_{nv}(y) = p * u(y_t(y)) + (1-p) * u(y) \tag{6}$$

$$y_t(y) = \max\{(1-\gamma_0)y, (1-\gamma_1)y - c_t\} \tag{7}$$

$$1 \geq \gamma_0 > \gamma_1 \geq 0 \tag{8}$$

To make the vaccination decision, the agent first decides whether to use medical care or not in the event of an infection. Because medical care is a binary choice, there is a threshold income $\tilde{y} = \frac{c_t}{\gamma_0 - \gamma_1}$ above which the agent uses medical care. As a result, the net income (i.e., consumption) in the event of an infection is given by:

$$y_t(y) = \begin{cases} (1-\gamma_0)y & \text{if } y < \tilde{y} \\ (1-\gamma_1)y - c_t & \text{if } y \geq \tilde{y} \end{cases} \tag{9}$$

Given the medical care decision, the agent decides whether to vaccinate or not by comparing the two expected utilities u_v and u_{nv}. Notice that this is essentially a choice between two gambles: one with the risk of side effects and the other with the risk of infection. This means that the degree of risk aversion is a key parameter in this decision. In what follows, I begin with the risk neutral case and later extend the discussion to the risk averse case. In this way, the role of risk aversion becomes clear.

3.1.1 Risk neutral case

Suppose the agent is risk neutral. The expected utilities with and without vaccination are given by:

$$u_v(y) = (1 - q\alpha)y - c_v \tag{10}$$

$$u_{nv}(y) = \begin{cases} (1-p\gamma_0)y & \text{if } y < \tilde{y} \\ (1-p\gamma_1)y - c_t & \text{if } y \geq \tilde{y} \end{cases} \tag{11}$$

To start, suppose the income of the agent is very low. By comparing the two expected utilities, it is immediately clear that the agent does not vaccinate if $y < \hat{y} = \frac{c_v}{1-q\alpha}$. In this income range, the expected utility from vaccination is always smaller than that from non-vaccination: $u_v(y) < 0 \leq u_{nv}(y)$. This is shown in Figure 7. This result comes from the fact that the agent has to pay the common cost c_v for vaccination while contracting the disease only involves time costs.

Figure 7: Risk neutral case

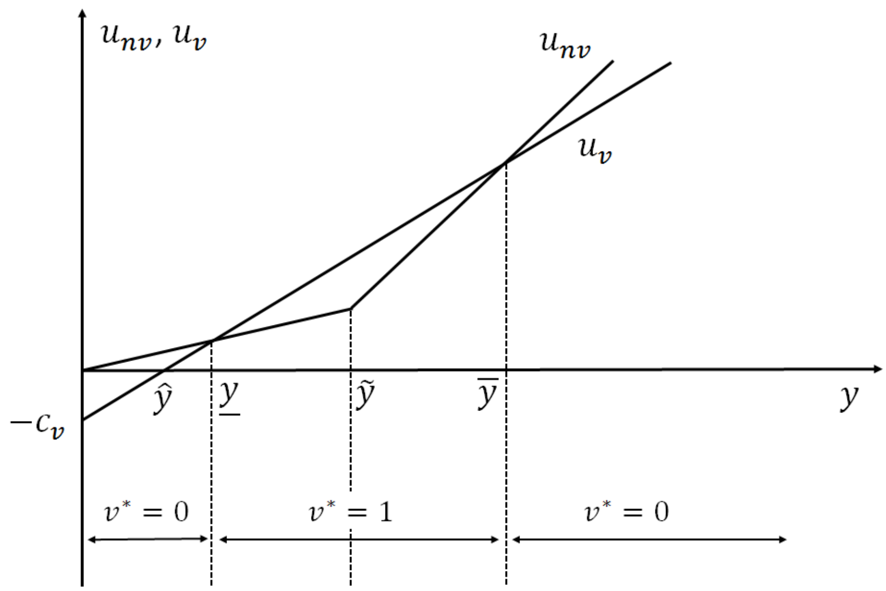

Now, suppose that the agent's income y is greater than $\hat{y}$. If the agent is to vaccinate at some level of income, the expected utility from vaccination must be larger than that from non-vaccination at that level. For this to be true, vaccination must be effective enough so that the marginal utility of income with vaccination $(1 - q\alpha)$ is larger than that without vaccination $(1 - p\gamma_0)$. Graphically, this means that the u_v curve is steeper than the u_{nv} curve, as in Figure 7. This condition reduces to:

$$p\gamma_0 > q\alpha \tag{12}$$

The left side of the inequality is the product of the risk of infection and the time loss when infected; the right side is the product of the risk of side effects and the time loss when side effects occur. Thus, the condition (12) states that the expected time loss from infection is larger than that from side effects. Under this condition, the agent starts using vaccination if the income exceeds $\underline{y} = \frac{c_v}{p\gamma_0 - q\alpha}$.

A caution is required, however, because if medical care is effective and inexpensive, the agent may always rely on medical care rather than vaccination. Graphically, this means that the u_{nv} curve always locates above the u_v curve so that these curves do not have any intersections. To exclude this possibility, it must be that the income at which the agent starts using vaccination $\underline{y}$ is lower than the income at which the agent starts using medical care $\tilde{y}$. This condition can be rewritten as:

$$\frac{p\gamma_0 - q\alpha}{c_v} > \frac{\gamma_0 - \gamma_1}{c_t} \tag{13}$$

In both sides of the inequality, the numerator is the effect of the intervention measured as the fraction of the work time saved, while the denominator is the cost of the intervention. Therefore, the condition (13) means that vaccination is more cost effective than medical care.

Finally, suppose the income of the agent is above $\tilde{y}$. For the agent not to vaccinate above a certain level of income, as in the low-income case, the expected utility from vaccination must be smaller than that from non-vaccination above the level, and this is only possible if medical care is effective enough. Graphically, this means the u_{nv} curve is steeper than the u_v curve in the high income range. A necessary condition for this is:

$$q\alpha > p\gamma_1 \tag{14}$$

This condition means that the expected time loss from side effects is greater than that from infection when medical care is utilized. Under this condition, the agent stops using vaccination once the income reaches $\bar{y} = \frac{c_t - c_v}{q\alpha - p\gamma_1}$ as shown in Figure 7.

For a given risk of infection p, the conditions (12), (13), and (14) ensure that the agent does not vaccinate if $0 \leq y < \underline{y}$, chooses to vaccinate if $\underline{y} < y \leq \bar{y}$, and does not vaccinate if $\bar{y} < y$. Recall however that p is endogenous. To ensure that these conditions are satisfied in equilibrium, therefore, we need to know p. One way to proceed is to introduce an assumption about the form of the income distribution to solve for p. Instead, here I keep the general income distribution and derive a set of sufficient conditions. Notice that p is bounded between $p(0)$ and $p(1)$, which are the risks of infection when the vaccination rate is 0 and 1, respectively. Therefore, a set of sufficient conditions that guarantee the hump-shaped vaccination decision rule is given by:

$$p(1)\gamma_0 > q\alpha > p(0)\gamma_1 \tag{15}$$

$$\frac{p(1)\gamma_0 - q\alpha}{c_v} > \frac{\gamma_0 - \gamma_1}{c_t} \tag{16}$$

Notice that conditions (15) and (16) do not hold if $p(1) = 0$. In general, however, $p(1) > 0$ either because there is a risk of importing the virus from other populations or because non-human reservoirs exist. The left inequality in condition (15) means that the minimum expected time loss due to infection (i.e., expected time loss when everyone else vaccinates) is larger than the expected time loss due to side effects, in the absence of medical care. The right inequality means that the maximum expected time loss due to infection is smaller than the expected time loss due to side effects, when medical care is utilized. Alternatively, condition (16) states that, even when everyone else vaccinates, vaccination is still more cost-effective than medical care. Note that if we assume a functional form for the income distribution, we will obtain less restrictive sufficient conditions.

3.1.2 Risk averse case

Suppose the agent is risk averse. To start, notice that just as in the risk neutral case, the agent does not vaccinate if $y < \hat{y}$ because $\forall y < \hat{y},\ u_v(y) < u(0) \leq u_{nv}(y)$.

Now, suppose the income of the agent is above $\hat{y}$. For the agent to vaccinate at some level of income, it is sufficient if the expected utility from vaccination is greater than that from non-vaccination at $y = \tilde{y}$, the income level at which the agent starts using medical care. This condition holds if:[24]

$$p\gamma_0 > \alpha \tag{17}$$
$$\frac{p\gamma_0 - \alpha}{c_v} > \frac{\gamma_0 - \gamma_1}{c_t} \tag{18}$$

Notice that condition (17) is a stricter version of (12) because it states that the expected time loss from infection is larger than the actual time loss from side effects. Correspondingly, condition (18) is a stricter version of (13) because it requires a higher cost-effectiveness for vaccination.

Finally, suppose the income of the agent is above $\tilde{y}$. For the agent to not vaccinate at higher levels of income, it is sufficient that there exists a threshold income $\acute{y}$ such that, for all the income above $\acute{y}$, the expected utility from vaccination is less than that from non-vaccination. This condition holds if:[25]

$$q\alpha > \gamma_1 \tag{19}$$

This condition is again a stricter version of (14) because it states that the expected time loss from side effects is larger than the actual time loss from infection when medical care is utilized.

For a given risk of infection p, conditions (17), (18), and (19) guarantee that there are two threshold income levels $\underline{y}$ and $\bar{y}$ at which the agent's vaccination decision flips. Recall however that p is endogenous and we need to solve for p to ensure that these conditions are satisfied in equilibrium. Without assuming any functional form for the income distribution, a set of sufficient conditions that guarantee the hump-shaped vaccination decision rule in equilibrium is given by:

$$p(1)\gamma_0 > \alpha, \quad q\alpha > \gamma_1 \tag{20}$$
$$\frac{p(1)\gamma_0 - \alpha}{c_v} > \frac{\gamma_0 - \gamma_1}{c_t} \tag{21}$$

The left inequality in condition (20) means that the minimum expected time loss due to infection is larger than the actual time loss due to side effects, in the absence of medical care. The right inequality means that the expected time loss due to side effects is larger than the actual time loss due to infection, when medical care is utilized. Correspondingly, condition (21) requires a higher cost effectiveness for vaccination. Notice again that these sufficient conditions are strong partly because I have not imposed any restriction on the

[24]Proof is as follows. On the one hand, $u_v = qu((1-\alpha)y - c_v) + (1-q)u(y - c_v) > u((1-\alpha)y - c_v)$. On the other hand, using Jensen's inequality, $u_{nv} = pu((1-\gamma_0)y) + (1-p)u(y) < u(p(1-\gamma_0) + (1-p))y) = u((1-p\gamma_0)y)$. Therefore, it is sufficient if $(1-\alpha)\tilde{y} - c_v > (1-p\gamma_0)\tilde{y}$. As $\tilde{y} = \frac{c_t}{\gamma_0 - \gamma_1}$, conditions (17) and (18) result.

[25]Proof is as follows. On the one hand, using Jensen's inequality, $u_v = qu((1-\alpha)y - c_v) + (1-q)u(y-c_v) < u(q((1-\alpha)y - c_v) + (1-q)(y-c_v)) = u((1-q\alpha)y - c_v)$. On the other hand, $u_{nv} = pu((1-\gamma_1)y - c_t) + (1-p)u(y) > u((1-\gamma_1)y - c_t)$. Therefore, it is sufficient if $\forall y > \acute{y}, (1-q\alpha)y - c_v < (1-\gamma_1)y - c_t$, which holds if $q\alpha > \gamma_1$.

income distribution. With a specific income distribution, these sufficient conditions will become less restrictive.

When compared to (15) and (16), the sufficient conditions (20) and (21) are stricter. In particular, the condition that makes the agent not vaccinate in the high-income range (i.e., $q\alpha > \gamma_1$) may be quite demanding. Notice, however, that this is not a necessary condition. As is clear from the proofs of the sufficient conditions, if the equilibrium vaccination rate is very small, this condition is overly strong. This is likely to occur when the disease is such that herd immunity is easy to achieve. Similarly, if the agent is less risk averse, the sufficient condition will approach to (15) so that condition (20) becomes overly strong as well. Therefore, both the agent's characteristics (i.e., the degree of risk aversion) and the disease's characteristics (i.e., the relationship between vaccination rates and the risk of infection) are important in determining whether the sufficient conditions are satisfied.

3.2 The vaccination bracket

For the rest of the paper, I assume that the sufficient conditions (20) and (21) hold. Let $G(y;p)$ be the difference between the two expected utilities:

$$G(y;p) = u_v(y) - u_{nv}(y;p) \tag{22}$$

where I explicitly denote the expected utility from non-vaccination as a function of p. When $G(y;p)$ is positive (negative), the agent does (does not) vaccinate. Under the sufficient conditions, $G(y;p)$ initially rises but then falls with income, having two intersections with the horizontal axis. Consequently, the agent's vaccination decision rule is given as a simple three-step function of income. There are two threshold incomes $\underline{y}$ and $\bar{y}$ where the agent's vaccination decision flips as shown in Figure 8. As a result, the agent neither vaccinates nor uses medical care if $0 \leq y \leq \underline{y}$; chooses to vaccinate if $\underline{y} < y \leq \bar{y}$; and does not vaccinate and relies on medical care in the event of an infection if $\bar{y} < y$. I name the income range between $\underline{y}$ and $\bar{y}$ as the vaccination bracket.

3.3 Equilibrium vaccination rate

The agent makes the vaccination decision based on the perception that the risk of infection p is given. However, the risk of infection is endogenously determined by the vaccination rate in the population. To determine the equilibrium vaccination rate, we need to aggregate the decision across agents. As each agent's vaccination decision depends on the vaccination rate and this in turn depends on these vaccination decisions, this is a fixed-point problem. To solve this problem, notice that Figure 8 relates individual agents' vaccination decisions to the vaccination rate in the population; the fraction of the population that falls into the vaccination bracket corresponds to the vaccination rate in the population.

In relating the vaccination bracket to the vaccination rate, notice that there are two cases. The first is that the maximum income of the population y_m is larger than $\bar{y}$ so that the

Figure 8: The vaccination bracket

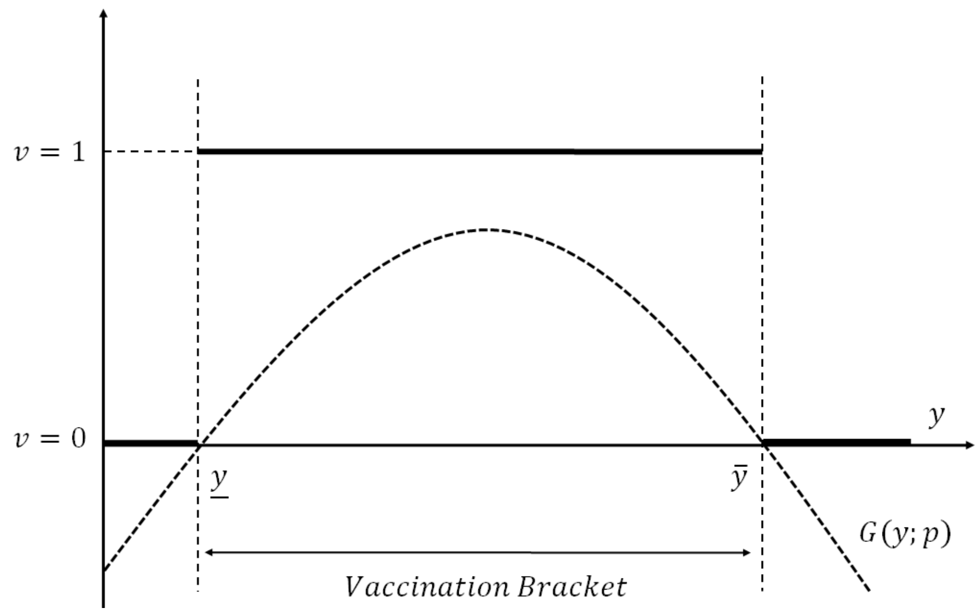

vaccination bracket exists strictly inside the income distribution.[26] I name this population as a developed population. In this case, the vaccination rate is given by $v = F(\bar{y}) - F(\underline{y})$, where F is the cumulative income distribution function. A developed population consists of three groups: low-income (i.e., $y < \underline{y}$) agents who do not vaccinate; middle-income ($\underline{y} < y \leq \bar{y}$) agents who vaccinate; and high-income (i.e., $y > \bar{y}$) agents who do not vaccinate. Therefore, a hump-shaped relationship between the vaccination rate and income emerges.

The second case is that the maximum income y_m is smaller than $\bar{y}$ so that the vaccination bracket does not fit inside the income distribution. I name this population as a developing population. In this case, the vaccination rate is given by $v = 1 - F(\underline{y})$. A developing population consists of two groups: low-income (i.e., $y < \underline{y}$) who do not vaccinate; and middle-to-high-income ($\underline{y} < y$) agents who vaccinate. Therefore, a positive relationship exists between the vaccination rate and income.

For now, let us focus on a developed population. Suppose that the risk of infection p increases. This will make vaccination more attractive and shift the $G(y; p)$ curve upward in Figure 8. As a result, the vaccination bracket expands and the vaccination rate in the population rises. Suppose instead that the risk of infection p decreases. This will make vaccination less attractive and shift the $G(y; p)$ curve downward. As a result, the vaccination bracket shrinks and the vaccination rate declines. Therefore, there is a positive relationship between the risk of infection p and the vaccination rate in the population.

Further, the risk of infection is determined by the vaccination rate: the higher the vaccination rate, the lower the risk of infection. Therefore, as the vaccination rate rises, the risk

[26]This occurs when if $y_m > k$ such that $G(k; p(v)) = 0, G(\underline{y}; p(v)) = 0, v = 1 - F(\underline{y})$. k is a threshold income at which the upper bound of the vaccination bracket coincides with the highest income in the population (i.e., $y_m = \bar{y}$).

Figure 9: The fixed-point problem

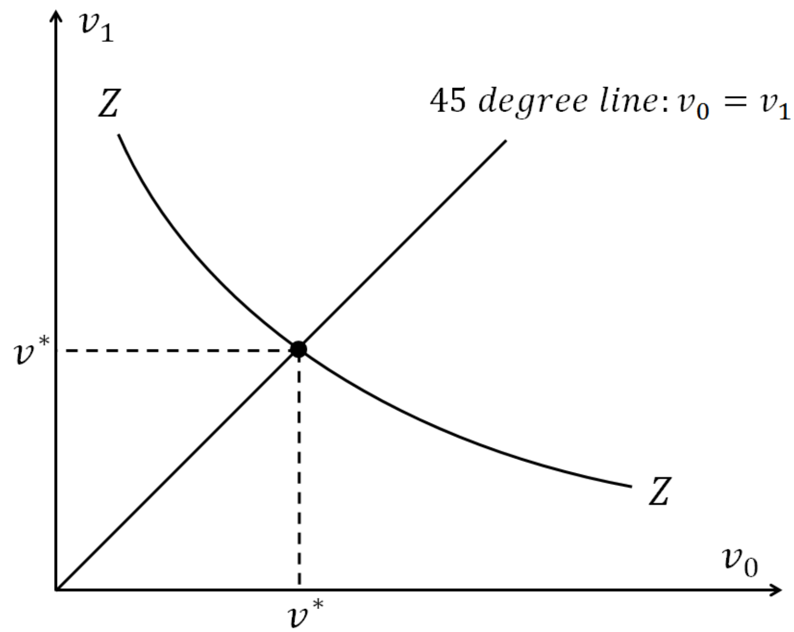

of infection p falls, and the resulting vaccination rate is lower. Therefore, there is a negative relationship between the initial vaccination rate v_0 and the resulting vaccination rate v_1. This is expressed as the ZZ curve in Figure 9. In an equilibirum, these two vaccination rates v_0 and v_1 must coincide, which is given by the intersection between the ZZ curve and the 45 degree line. The equilibrium vaccination rate is then given by v^*.

Two things are worth noting. First, the above argument is independent of income distribution. The negative relationship between the initial and resulting vaccination rates always holds, and the equilibrium vaccination rate is always uniquely determined for any income distribution. Second, the three-step vaccination decision rule holds in the equilibrium. The only difference here is that $G(y; p)$ is pinned down by the equilibrium vaccination rate v^* (i.e., $G(y; p(v^*))$). Therefore, in equilibrium, both low- and high-income agents do not vaccinate while middle-income agents choose to do so. Thus, the model replicates the hump shape found for many vaccines in Section 1. Using the same logic, in the case of developing populations, only low-income agents do not vaccinate while middle-to-high-income agents choose to do so in equilibrium. Therefore, the model also replicates the monotonic increase in vaccination rates observed for other vaccines in Section 1.

3.4 Connecting Theory to Empirics

Using the theory developed, we are now ready to see why the inclusion of region- and year-fixed effects was so critical in the empirical analysis. The model predicts a different equilibrium vaccination rate across regions or over time if underlying parameters are different, and these differences may mask the hump-shaped relationship between income and vaccination rates.

To see this, first suppose that we observe a population at different time periods, and that the cost of vaccination has fallen over time. This means that the vaccination bracket in Figure 8 expands over time, which shifts the ZZ curve upward in Figure 9. As a result, the equilibrium vaccination rate is higher in later periods. If the per capita income of this population is growing over the same period, we may observe a simple positive correlation between per capita income and vaccination rates. The same logic applies to the accumulation of knowledge about the importance of vaccination and the development of safer vaccines. Therefore, year-fixed effects are critical in the empirical analysis to control for this shift in the ZZ curve.

Suppose instead that we observe several populations with different levels of per capita income, and that the richer populations have a higher population density. Because a higher population density means a higher risk of infection, the vaccination bracket is larger for richer populations. This in turn means that the ZZ curve of a richer population locates above that of a poorer population, and as a result, the equilibrium vaccination rate is higher for richer populations. The same logic applies to the difference in institution and culture. Therefore, region-fixed effects are again critical in controlling for the difference in the locations of the ZZ curves.

4 Alternative explanations

In the previous section, I showed how effective medical care could be responsible for the hump-shaped relationship between income and childhood vaccination rates. This section considers other potential explanations. Any explanation must reconcile with the nonlinearity in vaccination decisions that an agent does not vaccinate when income is low, vaccinate when income is in the middle range, and does not vaccinate if income is high enough. In light of this, there are at least three alternative explanations: avoidance, differential risk, and preference. The three alternatives are all related to income but operate through different channels; while avoidance is similar in spirit to medical care, the differential risk is tied to the disease environment or information accessibility, and preference is tied to human nature.

Suppose there is no medical care, but instead there is a way to effectively protect a child from contracting a disease. This is an avoidance measure, and examples may include avoiding daycare and nurseries, undertaking home-schooling, and alternating children's diets to boost their immune systems. Suppose avoidance is a normal good,[27] is a binary choice, costs c_a, and the risk of infection can be reduced from p_0 to p_1 ($< p_0$) with this avoidance. Then, the

[27]The avoidance behaviour is an example of self-protective behaviour. The literature shows that self-protection may not be a normal good depending on the risk preference of the agent and the nature of the risk (e.g., Sweeney & Beard, 1992). It is beyond the scope of this paper to examine whether the avoidance measure is a normal good in the context of childhood vaccination.

agent's problem is given by:

$$\max\{u_v(y), u_{nv}(y)\}$$
$$u_v(y) = qu((1-\alpha)y - c_v) + (1-q)u(y-c_v) \tag{23}$$
$$u_{nv}(y) = \begin{cases} p_0 u((1-\gamma_0)y) + (1-p_0)u(y) & \text{if } y \leq \hat{y} \\ p_1 u((1-\gamma_0)y - c_a) + (1-p_1)u(y-c_a) & \text{if } y > \hat{y} \end{cases} \tag{24}$$

Notice that the structure of the problem is quite similar to the one analysed in the main section, and so is the argument. When the income is very low, the agent does not vaccinate due to the common cost of vaccination. As income increases, the agent starts using vaccination as long as it is effective. After income reaches a certain level, however, the agent starts using avoidance instead of vaccination if avoidance is effective enough. Therefore, with avoidance, we should be able to find a similar set of sufficient conditions that generates a hump shape.

In the set up above, I have assumed that the benefit of avoidance is given by the difference in probabilities, which in turn relies on the vaccination rate. In some cases, it may be reasonable to think that the gap between these two probabilities—the benefit of undertaking avoidance—is falling in the population's vaccination rate. This would imply, avoidance is most beneficial in an environment where more people are potentially sick and carriers of the disease. If this is true, then investments in avoidance work much like an investment in vaccination and I should be able to find a similar set of sufficient conditions that generates a hump-shaped relationship between vaccination decision and income. This set up might be reasonable for some types of avoidance behaviour such as wearing a mask to avoid germs, but it will not be appropriate for others.

In other cases, the marginal benefit of avoidance may rise in the population's vaccination rate. This implies avoidance is less beneficial in an environment where most people are potentially sick and carriers of the disease. For example, if not taking public transit is my avoidance but the population of workers at my office is highly contagious, then the benefit of this type of avoidance falls with greater prevalence. Therefore, this type of individual behaviour can lead to quite different outcomes than the previous type of avoidance. Agents can decide not to invest in avoidance if the population is expected to be quite contagious, and this in turn can produce the equilibrium result of more sickness. Alternatively, they may decide to avoid if the population is expected not to be contagious, and this in turn can produce the equilibrium result of less sickness because of less contact. This sort of expectations-driven behaviour is thought to have been important in the AIDS epidemic, where avoidance is modelled by limiting the number of sexual partners. In that case, an increase in the prevalence of AIDS can produce less avoidance by agents because its marginal benefit has fallen so dramatically. This possibility is known as fatalism in the literature and examined in detail by Kremer (1996) and Auld (2003, 2006). In the simple vaccination set up with medical care, I have developed that fatalism cannot exist, and the equilibrium is unique. In a more general environment with certain types of avoidance behaviour, multiple equilibria could result with some exhibiting fatalism. How this possibility affects my choice of empirical method and results is at present unclear. I hope to resolve this issue in future work.

Next, the hump shape may arise if high-income parents face a lower probability of their children contracting an infectious disease. This may be the case when a population is geographically segregated according to income, and the disease occurs mainly in poor neighbourhoods due to less sanitary environments. Similar to the previous cases, if income is low, the agent does not vaccinate because of the common cost of vaccination. Assuming vaccination is effective, the agent begins using vaccination as income rises. Once income reaches a certain level, however, the risk of infection drops greatly because the agent moves to a more sanitary area. Consequently, the agent stops using vaccination because the risk now is negligible. Thus, a hump-shaped relationship between income and vaccination decision emerges.

Notice that the discussion above does not necessarily require that the actual risk of infection varies by income; it suffices if the "perceived" risk of infection is smaller for high-income agents. Once this is understood, it is also clear that the hump shape may occur if the perceived risk of side effects is larger for high-income agents. These differential perceptions across income levels may arise due to the different information that parents have access to. For example, high-income agents can conduct more extensive research about the potential side effects of vaccination, and may think that the risk of side effects is higher than low-income agents think it is.

Finally, it is possible that the hump shape is simply driven by preference: vaccination is a normal good when income is low but it becomes an inferior good when income reaches a high enough level. This may be the case if agents prefer a "natural life" and vaccination is considered "unnatural". If this is the case, vaccination involves a trade-off between the protection from illness and the disutility from unnatural choice. Again, when income is low, the agent does not vaccinate due to the common cost of vaccination. As long as vaccination is effective, the agent begins using vaccination as income rises. When income reaches a certain level, however, the agent stops using vaccination because the disutility from vaccination is so large that it is better for the agent to use its income for consumption. Thus, the agent's vaccination decision rule forms a hump-shaped relationship with income.

5 Conclusion

The occasional outbreak of various diseases remains a serious public health concern in even the most developed countries. The main cause of these outbreaks is that vaccination rates are not high enough to achieve herd immunity. Traditionally, low vaccination rates were thought to be a result of low income people having limited access to medical facilities. While this is surely true, recently, there is a growing recognition that certain high income groups deliberately choose not to vaccinate their children. The aim of this paper was to examine whether, and if so, why, there is systematic tendency for high-income parents to choose not to vaccinate their children.

I started my investigation with country-level data. While the raw data shows a simple positive correlation between per capita income and vaccination rates, this relationship changes dramatically once I control for region- and year-fixed effects: Vaccination rates exhibit a hump-shaped relationship with income. This result is significant both statistically and economically. Take polio vaccine as an example. The result suggests that the polio vaccination rate peaks at GDP per capita of $28,000; many developed countries have already passed this income level. Assuming an annual growth rate of 2 %, the model predicts that a country with per capita income of $5,000 will experience 2.1% increase in the polio vaccination rate during the next 20 years; while a country with per capita income of $30,000 will experience 4.7% drop in the polio vaccination rate during the same period.

Because the country-level analysis may suffer from omitted variables, reverse causality, and measurement errors, I moved on to county-level data in the U.S., which has much less concern for these potential issues. There, I found a similar pattern. For example, the result suggests that the MMR vaccination rate peaks at per capita income of $35,000; when per capita income increases by $1,000, the vaccination rate increases by 0.17 percentage points when per capita income is $20,000, but declines by 0.1 percentage points when per capita income is $50,000. To further mitigate the concern for omitted variable bias, I proceeded to individual-level data in the U.S. that allows me to control for state-year-fixed effects and detailed family characteristics. The result was once again similar. For instance, DTP4 vaccination rate peaks at per capita income of $25,000; when per capita income rises by $1,000, the probability of a child being up-to-date rises by 0.24 percentage points if per capita income is $15,000, but decreases by 0.28 percentage points if per capita income is $30,000.

Although each dataset has its own limitations, these results collectively point to the existence of a phenomenon that vaccination rates initially rise but then fall as income increases. I name this phenomenon as "Vaccination Kuznets Curve (VKC)." This empirical results then raised a question: What is driving this hump shape? One possibility is the prevalence-elastic prevention demand proposed by Geoffard & Philipson (1997), which states that the demand for prevention decreases as disease prevalence becomes lower. This theory may explain why vaccination rates are low in high-income countries. It does not, however, tell us what generates the hump shape, and more importantly, it cannot be reconciled with the county- and individual-level results in the U.S.

To examine this puzzle, I built a model of parents' vaccination decisions. The model explains that low-income parents do not vaccinate because the common cost of vaccination, such as transportation to a clinic, is an obstacle for them. In contrast, high-income parents do not vaccinate due to the opportunity cost of the potential side effects of vaccination. They prefer taking the risk of infection when medical care (is believed to) effectively mitigates the opportunity cost of infection. A similar trade-off exists when parents use avoidance measures to reduce the risk of infection instead of using medical care to mitigate disease symptoms.

The policy implications from the model are straightforward. To encourage more low-income parents to vaccinate, it is important to further reduce the common cost of vaccination. To encourage more high-income parents to vaccinate, it is important to continue our quest for safer vaccines while disseminating more accurate information about the effectiveness of medical care and vaccination. If parents overestimate the effect of alternative measures to protect their children, they are less likely to vaccinate their children. This possibility seems to be underemphasised in policy circles.[28] The model developed in this paper is useful because it helps us identify potential channels through which public health policies can effectively raise the vaccination rate in our society.

There are of course other potential explanations for the VKC, and policy prescriptions will be different depending on the explanation. I do not even exclude the possibility that it is not income itself but other region-specific time-varying factors that generate the hump shape. Whatever the mechanism is, however, the VKC indicates the possibility that vaccination rates will fall and disease outbreaks will eventually rise in the near future. The raw data has not shown a clear decline in vaccination rates partly because public health authorities have been successful in raising the public's interest in vaccination, but this trend is not guaranteed to continue. In fact, recent outbreaks of childhood diseases such as measles and pertussis in North America and Europe may be just the beginning of a scenario suggested by this paper.

[28]Reich (2014) uses an interview survey and finds that the mothers who do not vaccinate their children think they can reduce the risk of disease with avoidance measures.

6 Appendix

6.1 Country-level results: Cubic regression

Table 7: Estimation results, Cross country

	Pol3	DTP3	BCG	MCV	Hib3	HepB3
Income	14.240	7.861	10.108	14.231	31.654^b	13.966
	(8.560)	(8.530)	(12.883)	(9.745)	(14.843)	(19.913)
Income squared	-3.165	-0.787	-1.979	-3.202	-8.209^b	-3.794
	(1.966)	(2.363)	(3.235)	(2.392)	(3.392)	(4.525)
Income cubed	0.149	-0.047	-0.037	0.186	0.531^b	0.283
	(0.153)	(0.200)	(0.236)	(0.182)	(0.229)	(0.309)
Population	0.046	0.144	0.538^a	0.018	-0.323^b	0.673
	(0.174)	(0.200)	(0.139)	(0.179)	(0.139)	(0.436)
Population density	-1.839	-3.284^c	-0.938	-3.330^b	-6.252	-11.746^a
	(1.576)	(1.899)	(2.103)	(1.569)	(9.588)	(3.986)
Share of pop: 15-64	0.120	0.458	0.128	0.221	1.462^c	-0.483
	(0.357)	(0.467)	(0.432)	(0.432)	(0.855)	(1.189)
Share of pop: >64	-0.235	-0.022	0.933	1.380^c	2.584	4.884
	(0.827)	(0.666)	(0.901)	(0.752)	(1.618)	(3.010)
Share of pop: female	-0.653	-0.398	-0.446	-1.607^c	1.332	-0.600
	(0.769)	(0.793)	(0.559)	(0.816)	(1.104)	(1.123)
Share of pop: rural	-0.345	-0.187	0.160	0.158	1.543^b	0.912
	(0.209)	(0.210)	(0.189)	(0.248)	(0.592)	(0.671)
N	2034	2030	1154	1990	964	938
Countries	65	65	42	65	63	56
P(G=0)	0.003	0.053	0.000	0.031	0.042	0.809
R2	0.400	0.488	0.506	0.586	0.309	0.377
RMSE	8.232	8.027	7.705	9.074	9.169	12.004

Notes: Income is in $10,000. Standard errors are clustered at the country level. a, b, and c mean statistical significance at the 1 %, 5%, and 10% levels, respectively. All the specifications include country- and year-fixed effects. P(G=0) shows p-values for the joint test of income variables.

6.2 Country-level results: Dummy regression

Table 8: Estimation results, Cross country

	Pol3	DTP3	BCG	MCV	Hib3	HepB3
Per capita GDP 1,000-1,999	3.093	4.544	6.438[b]	3.804	19.868[a]	5.692
	(3.244)	(3.272)	(2.802)	(2.942)	(6.883)	(8.618)
Per capita GDP 2,000-2,999	5.611	10.201[b]	3.867	3.843	20.419[a]	6.095
	(4.592)	(4.729)	(9.501)	(4.768)	(7.592)	(11.707)
Per capita GDP 3,000-3,999	3.176	10.352	1.268	4.352	16.579[b]	0.513
	(5.372)	(6.572)	(9.844)	(6.036)	(7.898)	(12.224)
Per capita GDP 4,000-4,999	-0.934	7.647	2.449	1.605	14.645[c]	0.762
	(5.411)	(8.951)	(12.065)	(7.020)	(7.963)	(12.801)
Per capita GDP 5,000-5,999	-4.738	6.944	-37.152[a]	-0.853	4.938	
	(6.052)	(10.850)	(10.799)	(8.502)	(9.296)	
Per capita GDP 6,000-6,999	-6.921	2.287	-43.371[a]	-3.468	2.344	2.533
	(5.840)	(9.649)	(10.434)	(7.892)	(9.150)	(3.405)
Per capita GDP 7,000-7,999	-33.684[a]	-25.530[a]	-57.481[a]	-14.380[c]	-1.078	-0.737
	(5.972)	(8.616)	(10.222)	(7.306)	(0.992)	(2.269)
Per capita GDP 8,000-8,999	-41.372[a]	-30.130[a]	-64.432[a]	-15.373[b]		
	(6.480)	(9.057)	(10.513)	(7.345)		
Population	0.053	0.170	0.550[a]	0.008	-0.262[b]	0.636
	(0.172)	(0.191)	(0.148)	(0.175)	(0.110)	(0.490)
Population density	-1.308	-3.281[b]	0.742	-2.432	-3.351	-10.241[b]
	(1.294)	(1.525)	(2.519)	(1.568)	(6.664)	(5.003)
Share of pop: 15-64	0.307	0.456	0.473	0.446	1.972[b]	-0.381
	(0.332)	(0.393)	(0.425)	(0.371)	(0.797)	(1.244)
Share of pop: >64	0.255	0.167	1.381	1.759[b]	3.317[b]	5.140[c]
	(0.880)	(0.545)	(0.948)	(0.753)	(1.574)	(3.065)
Share of pop: female	-0.435	-0.344	-0.004	-1.308[c]	1.835[c]	-0.344
	(0.731)	(0.720)	(0.456)	(0.767)	(0.937)	(1.198)
Share of pop: rural	-0.356[c]	-0.211	0.086	0.139	1.278[b]	0.814
	(0.213)	(0.206)	(0.188)	(0.254)	(0.568)	(0.661)
N	2034	2030	1154	1990	964	938
Countries	65	65	42	65	63	56
P(G=0)	0.000	0.000	0.000	0.003	0.015	0.938
R2	0.401	0.498	0.551	0.584	0.338	0.379
RMSE	8.240	7.957	7.363	9.110	8.992	12.010

Notes: Income is in $10,000. Standard errors are clustered at the country level. a, b, and c mean statistical significance at the 1 %, 5%, and 10% levels, respectively. All the specifications include country- and year-fixed effects. P(G=0) shows p-values for the joint test of income variables.

6.3 County-level results

Table 9: Estimation results, US county

	Pol3	DTP4	MMR1	Hib3	HepB3	VRC1	PCV4
Income	-0.538	-0.018	5.000[a]	-5.056[c]	11.432[a]	21.726[a]	46.622[a]
	(2.386)	(3.610)	(1.851)	(2.742)	(3.688)	(7.706)	(16.645)
Income squared	0.056	0.090	-0.981[a]	0.790[c]	-2.203[a]	-3.987[a]	-7.962[a]
	(0.399)	(0.530)	(0.310)	(0.470)	(0.614)	(1.252)	(2.747)
Income cubed	-0.001	-0.000	0.050[a]	-0.036	0.108[a]	0.187[a]	0.382[a]
	(0.021)	(0.025)	(0.016)	(0.025)	(0.031)	(0.061)	(0.132)
Population	-0.864	3.966[a]	0.918	1.453	0.603	-9.656[b]	27.872[a]
	(1.362)	(1.467)	(0.920)	(1.294)	(1.830)	(3.731)	(8.996)
Population density	-3.321	-2.872	-1.678	0.675	-5.764[a]	4.966	7.832
	(2.136)	(2.378)	(1.594)	(1.874)	(1.390)	(3.215)	(19.337)
Share of pop: 15-64	-13.985	11.363	-1.499	19.916	-42.932[c]	-81.668	161.632
	(15.072)	(22.518)	(12.369)	(15.512)	(25.343)	(55.128)	(139.589)
Share of pop: >64	-21.456	-21.920	-4.094	-34.571[c]	38.313	-23.930	174.772
	(18.693)	(25.008)	(15.265)	(18.984)	(29.241)	(67.089)	(166.558)
Share of pop: female	3.371	49.533[c]	-4.280	31.076	-115.876[a]	-200.079[b]	82.867
	(19.957)	(29.350)	(19.872)	(22.863)	(34.637)	(86.352)	(353.248)
Share of pop: white	9.531	-14.821	7.168	-1.095	57.666[a]	53.134[c]	211.134[b]
	(12.605)	(15.424)	(7.352)	(10.869)	(15.707)	(32.085)	(98.782)
Share of pop: black	13.964	-12.000	6.219	-8.612	61.881[a]	13.082	211.520
	(14.420)	(17.616)	(9.584)	(11.827)	(16.934)	(37.296)	(128.936)
Share of pop: native	39.956	-38.589	36.014	-3.833	33.267	-150.023	104.641
	(41.208)	(73.068)	(31.675)	(66.210)	(62.685)	(136.327)	(192.652)
N	1287	1287	1287	1287	1287	1099	556
Counties	229	229	229	229	229	229	229
P(G=0)	0.916	0.046	0.004	0.114	0.000	0.000	0.040
R2	0.470	0.353	0.236	0.328	0.847	0.958	0.942
RMSE	1.536	2.251	1.367	1.380	2.436	3.966	3.633

Notes: Income is in $10,000. Standard errors are clustered at the county level. a, b, and c mean statistical significance at the 1 %, 5%, and 10% levels, respectively. All the specifications include county- and year-fixed effects. P(G=0) shows p-values for the joint test of income variables.

6.4 Individual-level results

Table 10: Vaccine results, US individual, linear regression

	Pol3	DTP4	MCV1	HiB3	HepB3	VRC1	PCV4	ROT3
Income	0.141	2.553	0.333	2.997^b	0.271	-1.044	6.271^a	13.833^a
	(1.363)	(1.877)	(1.356)	(1.506)	(1.393)	(1.462)	(2.103)	(3.169)
Income squared	0.316	0.819	0.476	-0.477	0.922	1.673	-0.422	-5.849^b
	(1.089)	(1.478)	(1.084)	(1.229)	(1.141)	(1.168)	(1.696)	(2.717)
Income cubed	-0.037	-0.382	-0.114	-0.108	-0.280	-0.340	-0.248	1.029
	(0.256)	(0.341)	(0.255)	(0.297)	(0.278)	(0.275)	(0.404)	(0.671)
Child age: 24-29 month	1.688^a	7.731^a	2.911^a	1.399^a	0.984^a	2.801^a	1.796^a	-2.535^a
	(0.309)	(0.449)	(0.316)	(0.319)	(0.323)	(0.352)	(0.469)	(0.727)
Child age: 30-35 month	2.472^a	10.705^a	3.918^a	2.064^a	1.953^a	3.764^a	2.113^a	-7.275^a
	(0.299)	(0.430)	(0.301)	(0.315)	(0.308)	(0.338)	(0.461)	(0.705)
N of HH member: 3	0.463	-0.713	-0.258	-0.295	0.100	0.151	-1.104	3.029
	(0.911)	(1.202)	(0.885)	(0.934)	(0.920)	(1.101)	(1.372)	(2.037)
N of HH member: 4	0.379	-1.048	-0.361	-0.229	0.606	0.497	-0.788	2.743
	(1.016)	(1.361)	(0.997)	(1.044)	(1.023)	(1.208)	(1.535)	(2.278)
N of HH member: 5	-1.204	-3.027^b	-1.445	-1.218	-0.621	-0.876	-2.143	-0.046
	(1.032)	(1.381)	(1.011)	(1.062)	(1.044)	(1.220)	(1.558)	(2.317)
N of HH member: 6	0.061	-1.326	-1.567	0.007	0.193	-1.259	-1.355	3.658
	(1.142)	(1.545)	(1.108)	(1.164)	(1.165)	(1.334)	(1.758)	(2.632)
N of HH member: 7	-1.680	-3.869^b	-2.686^b	-1.844	-1.198	-2.303	-2.632	2.299
	(1.284)	(1.727)	(1.219)	(1.355)	(1.323)	(1.455)	(1.921)	(2.809)
N of HH member: 8+	-2.694^c	-4.573^b	-4.195^a	-2.582^c	-2.615^c	-4.429^a	-5.521^a	-0.125
	(1.402)	(1.834)	(1.336)	(1.450)	(1.418)	(1.550)	(2.040)	(2.981)
Respondent: mother	0.588	11.311	-3.459	9.331	-1.808	-4.856	7.377	37.971^b
	(8.028)	(27.565)	(8.547)	(17.968)	(8.294)	(8.614)	(25.602)	(18.804)
Respondent: father	-0.219	9.721	-4.698	7.989	-2.833	-5.512	5.160	36.159^c
	(8.032)	(27.568)	(8.553)	(17.970)	(8.300)	(8.621)	(25.607)	(18.818)
Respondent: grandparent	0.404	10.665	-2.679	8.836	-2.979	-3.981	5.609	33.852^c
	(8.066)	(27.583)	(8.570)	(17.986)	(8.335)	(8.642)	(25.625)	(18.855)
Respondent: other	-1.914	10.932	-2.514	5.659	-2.904	-2.196	4.224	33.996^c
	(8.212)	(27.645)	(8.675)	(18.063)	(8.457)	(8.744)	(25.704)	(19.072)
Respondent: dont't know	-5.522	13.530	-8.734	4.067	9.263	-8.001	11.737	37.376
	(14.199)	(29.395)	(14.471)	(21.369)	(8.512)	(14.820)	(27.467)	(23.045)
N of children: 2-3	-0.656	-1.186	-0.165	-0.945^c	-0.972^c	-0.644	-1.039	0.324
	(0.502)	(0.737)	(0.516)	(0.531)	(0.527)	(0.562)	(0.816)	(1.267)
N of children: 4+	-3.571^a	-6.678^a	-2.035^b	-4.292^a	-3.138^a	-2.668^a	-6.090^a	-8.506^a
	(0.848)	(1.192)	(0.803)	(0.893)	(0.877)	(0.897)	(1.294)	(1.936)
Mother's education: =12	1.475^a	2.651^a	1.386^a	1.751^a	1.309^a	1.558^a	3.586^a	4.151^a
	(0.498)	(0.681)	(0.479)	(0.524)	(0.497)	(0.525)	(0.738)	(1.092)
Mother's education: >12, non college	2.633^a	3.996^a	1.568^a	2.459^a	1.892^a	2.185^a	4.938^a	5.923^a
	(0.529)	(0.713)	(0.532)	(0.549)	(0.536)	(0.575)	(0.777)	(1.150)
Mother's education: college	2.864^a	6.497^a	2.623^a	3.215^a	1.393^b	2.850^a	7.085^a	9.388^a
	(0.554)	(0.728)	(0.543)	(0.560)	(0.559)	(0.593)	(0.807)	(1.210)
Firstborn: Yes	0.387	1.879^a	1.028^a	0.523	-0.007	0.594	2.343^a	3.336^a
	(0.335)	(0.468)	(0.321)	(0.348)	(0.344)	(0.366)	(0.521)	(0.849)
Hispanic: Yes	0.269	-8.664	-12.101	-13.672	2.628	-0.013	1.721	1.758
	(3.541)	(13.954)	(13.375)	(12.695)	(2.438)	(17.334)	(19.775)	(2.252)
Language: English	-1.279	-1.985	-1.093	1.769	-0.901	-1.083	3.230^c	3.589
	(1.105)	(1.585)	(1.047)	(1.301)	(1.057)	(1.273)	(1.789)	(2.488)
Language: Spanish	2.871^b	5.826^a	3.625^a	6.183^a	2.613^b	3.983^a	11.090^a	15.721^a
	(1.232)	(1.783)	(1.176)	(1.433)	(1.215)	(1.423)	(2.004)	(2.844)
Mother's age: 20-29	-0.086	-2.868^c	-1.721^c	-0.800	0.593	-1.579	-4.356^a	-1.395
	(1.174)	(1.468)	(0.940)	(1.192)	(1.207)	(1.025)	(1.568)	(2.373)
Mother's age: >=30	0.765	0.124	-0.676	0.388	0.740	-0.455	-1.004	2.374
	(1.192)	(1.500)	(0.966)	(1.217)	(1.234)	(1.053)	(1.599)	(2.423)
Married	0.398	1.063^b	0.466	0.529	-0.236	-0.055	0.876	-0.479
	(0.381)	(0.539)	(0.364)	(0.415)	(0.391)	(0.400)	(0.577)	(0.854)
Moved	-4.948^a	-8.339^a	-4.915^a	-5.954^a	-4.398^a	-4.814^a	-11.103^a	-14.369^a
	(0.579)	(0.742)	(0.589)	(0.643)	(0.580)	(0.604)	(0.777)	(1.087)
Race: White only	0.404	1.126	0.444	-0.237	-0.985	-0.087	1.861	0.687
	(0.939)	(1.313)	(0.892)	(0.875)	(0.863)	(0.923)	(1.499)	(2.191)
Race: Blach only	-1.886	-3.468^c	0.916	-1.709	-1.885	-0.476	-3.624	-2.652
	(1.502)	(2.068)	(1.302)	(1.350)	(1.404)	(1.421)	(2.253)	(3.150)
Race: Hispanic	-0.307	7.336	10.674	15.279	-2.522	-1.030	-2.198	
	(3.582)	(13.975)	(13.385)	(12.706)	(2.490)	(17.342)	(19.793)	
Race: Non-Hispanic, white only	-0.904	-1.924	-2.212^b	0.406	-0.793	-2.678^a	-0.993	0.041
	(1.018)	(1.435)	(0.968)	(0.973)	(0.943)	(1.014)	(1.625)	(2.379)
Race: Non-Hispanic, black only	1.205	1.529	-1.779	1.574	0.467	-0.729	1.938	-0.675
	(1.566)	(2.182)	(1.374)	(1.450)	(1.474)	(1.510)	(2.396)	(3.369)
Female	0.269	0.237	0.186	-0.059	-0.022	0.054	0.267	0.640
	(0.236)	(0.335)	(0.237)	(0.250)	(0.247)	(0.265)	(0.368)	(0.569)
N	161002	161002	161002	161002	161002	161002	161002	92501
P(G=0)	0.044	0.000	0.014	0.000	0.005	0.000	0.000	0.000
R2	0.024	0.055	0.024	0.045	0.020	0.025	0.096	0.107

Notes: Income is in \$10,000. a, b, and c mean statistical significance at the 1 %, 5%, and 10% levels, respectively. All the specifications include state-year-fixed effects and other control variables. P(G=0) shows p-values for the joint test of income variables.

References

Andreoni, J. & Levinson, A. (2001). 'The simple analytics of the environmental kuznets curve.' *Journal of Public Economics*, vol. 80(2), pp. 269–286.

Auld, M. C. (2003). 'Choices, beliefs, and infectious disease dynamics', *Journal of Health Economics*, vol. 22(3), pp. 361–377.

Auld, M. C. (2006). 'Estimating behavioral response to the AIDS epidemic', *Contributions in Economic Analysis & Policy*, vol. 5(1), pp. 1–29.

Barrett, S. & Hoel, M. (2007). 'Optimal disease eradication', *Environment and Development Economics*, vol. 12(5), pp. 627–652.

Brito, D. L., Sheshinski, E., & Intriligator, M. D. (1991). 'Externalities and compulsary vaccinations', *Journal of Public Economics*, vol. 45(1), pp. 69–90.

Brock, W. A. & Taylor, M. S. (2010). 'The green solow model', *Journal of Economic Growth*, vol. 15(2), pp. 127–153.

Burton, A., Monasch, R., Lautenbach, B., Gacic-Dobo, M., Neill, M., Karimov, R., Wolfson, L., Jones, G., & Birmingham, M. (2009). 'WHO and UNICEF estimates of national infant immunization coverage: methods and processes', *Bulletin of the World Health Organization*, vol. 87, pp. 535–541.

Chen, F. & Toxvaerd, F. (2014). 'The economics of vaccination', *Journal of Theoretical Biology*, vol. 363, pp. 105–117.

Cole, M. A., Rayner, A. J., & Bates, J. M. (1997). 'The environmental kuznets curve: an empirical analysis', *Environment and Development Economics*, vol. 2(4), pp. 401–416.

Copeland, B. R. & Taylor, M. S. (2003). *Trade and the Dnvironment: Theory and Dvidence.* Princeton University Press.

Francis, P. J. (1997). 'Dynamic epidemiology and the market for vaccinations', *Journal of Public Economics*, vol. 63(3), pp. 383–406.

Fry, S. J. (2011). 'Barriers to up-to-date pertussis immunization in Oregon children', *Scholar Archive*. Paper 594.

Garrett, L. (1994). *The coming plague: newly emerging diseases in a world out of balance.* New York: Farrar, Straus and Giroux.

Geoffard, P.-Y. & Philipson, T. (1997). 'Disease eradication: private versus public vaccination', *The American Economic Review*, vol. 87(1), pp. 222–230.

Gersovitz, M. (2003). 'Births, Recoveries, Vaccinations, and Externalities', in (R.J. Arnott *et al*, eds.), *Economics for an imperfect world: Essays in honor of Joseph E. Stiglitz*, pp. 469–483. MIT Press.

Gersovitz, M. & Hammer, J. S. (2003). 'Infectious diseases, public policy, and the marriage of economics and epidemiology', *The World Bank Research Observer*, vol. 18(2), pp. 129–157.

Gersovitz, M. & Hammer, J. S. (2004). 'The economical control of infectious diseases', *The Economic Journal*, vol. 114(492), pp. 1–27.

Goldman, S. M. & Lightwood, J. (2002). 'Cost optimization in the SIS model of infectious disease with treatment', *Topics in Economic Analysis & Policy*, vol. 2(1), pp. 1–22.

Grossman, G. M. & Krueger, A. B. (1991). 'Environmental impacts of a North American Free Trade Agreement', *NBER Working Paper No. 3914*.

Grossman, G. M. & Krueger, A. B. (1994). 'Economic Growth and the Environment', *The Quarterly Journal of Economics*, vol. 110(2), pp. 353–377.

Harbaugh, W. T., Levinson, A., & Wilson, D. M. (2000). 'Reexamining the empirical evidence for an environmental Kuznets curve', *Review of Economics and Statistics*, vol. 84, pp. 541–551.

John, A. & Pecchenino, R. (1994). 'An overlapping generations model of growth and the environment', *The Economic Journal*, vol. 104(427), pp. 1393–1410.

Klevens, R. M. & Luman, E. T. (2001). 'US children living in and near poverty: risk of vaccine-preventable diseases', *American Journal of Preventive Medicine*, vol. 20(4), pp. 41–46.

Kremer, M. (1996). 'Integrating behavioral choice into epidemiological models of aids', *The Quarterly Journal of Economics*, vol. 111(2), pp. 549–573.

Kureishi, W. (2009). 'Partial vaccination programs and the eradication of infectious diseases', *Economics Bulletin*, vol. 29(4), pp. 2758–2769.

List, J. A. & Gallet, C. A. (1999). 'The environmental kuznets curve: does one size fit all?', *Ecological Economics*, vol. 31(3), pp. 409–423.

Lopez, R. (1994). 'The environment as a factor of production: the effects of economic growth and trade liberalization', *Journal of Environmental Economics and Management*, vol. 27(2), pp. 163–184.

Reich, J. A. (2014). 'Neoliberal mothering and vaccine refusal imagined gated communities and the privilege of choice', *Gender & Society*, vol. 28(5), pp.679-704.

Rowthorn, B. R. & Toxvaerd, F. (2012). 'The optimal control of infectious diseases via prevention and treatment', *CEPR Discussion Paper No. DP8925*.

Shafik, N. & Bandyopadhyay, S. (1991). 'Economic growth and environmental quality: time-series and cross-country evidence', Background Paper for the World Development Report 1992, The World Bank, Washington, DC.

Smith, J. P., Humiston, G. S., Marcuse, K. E., Zhao, Z., Dorell, G. C., Howes, C., & Hibbs, B. (2011). 'Parental delay or refusal of vaccine doses, childhood vaccination coverage at 24 months of age, and the health belief model', *Public Health Reports*, vol. 126, Suppl 2, pp. 135–146.

Smith, J. P. & Singleton, A. J. (2011). 'County-level Trends in Vaccination Coverage Among Children Aged 19-35 Months: United States, 1995-2008', *MMWR*, vol. 60(4), pp.1–86.

Smith, P. J., Chu, S. Y., & Barker, L. E. (2004). 'Children who have received no vaccines: who are they and where do they live?', *Pediatrics*, vol. 114(1), pp. 187–195.

Stokey, N. L. (1998). 'Are there limits to growth?', *International Economic Review*, vol. 39(1), pp. 1–31.

Sweeney, G. H. & Beard, T. R. (1992). 'The comparative statistics of self-protection', *Journal of Risk and Insurance*, vol. 59(2), pp. 301–309.

Toxvaerd, F. (2010a). 'Infection, acquired immunity and externalities in treatment', *CEPR Discussion Paper No. DP8111.*

Toxvaerd, F. (2010b). 'Recurrent infection and externalities in prevention', *CEPR Discussion Paper No. DP8112.*

Troesken, W. (2015). *The Pox of Liberty: How the Constitution Left Americans Rich, Free, and Prone to Infection*. University of Chicago Press.

Wei, F., Mullooly, J. P., Goodman, M., McCarty, M. C., Hanson, A. M., Crane, B., & Nordin, J. D. (2009). 'Identification and characteristics of vaccine refusers', *BMC pediatrics*, vol. 9(18).

Wiemer, C. (1987). 'Optimal disease control through combined use of preventive and curative measures', *Journal of Development Economics*, vol. 25(2), pp. 301–319.

Wu, A. C., Wisler-Sher, D. J., Griswold, K., Colson, E., Shapiro, E. D., Holmboe, E. S., & Benin, A. L. (2008). 'Postpartum mothers' attitudes, knowledge, and trust regarding vaccination', *Maternal and Child Health Journal*, vol. 12(6), pp. 766–773.

Xu, X. (1999). 'Technological improvements in vaccine efficacy and individual incentive to vaccinate', *Economics Letters*, vol. 65(3), pp. 359–364.

Yang, Y. T., Delamater, P. L., Leslie, T. F., & Mello, M. M. (2016). 'Sociodemographic predictors of vaccination exemptions on the basis of personal belief in California', *American Journal of Public Health*, vol. 106(1), pp. 172–177.