

Measuring Social and Externality Benefits of Influenza Vaccination *

Corey White

California Polytechnic State University, San Luis Obispo
cwhite46@calpoly.edu

This Version: June 22, 2018

Abstract

Vaccination represents a canonical example of externalities in economics, yet there are few estimates of their magnitudes. I provide evidence on the social and externality benefits of influenza vaccination in two settings. First, using pre-existing differences in state-level vaccination rates interacted with exogenous annual variation in vaccine quality, I estimate the impacts of aggregate vaccination rates on mortality and work absences in the United States. Scaled nationally, I find that a one percentage point increase in the vaccination rate results in approximately 800 fewer deaths and 15 million fewer work hours lost due to illness each year. The mortality reductions are concentrated among individuals 75 and older, but over half of the effect is attributable to the vaccination of people under 75, suggesting a considerable externality effect. Second, I examine a setting in which vaccination is targeted at a group with extremely high externality benefits: vaccination mandates for health care workers. I find that mandates lead to reductions in hospital diagnoses for influenza in affected counties, consistent with substantial externality impacts. For both the general population and the population of health care workers, the estimates suggest that programs increasing vaccine take-up are likely to be cost-effective under reasonable assumptions about the costs.

JEL Codes: I12, I18, D62, H23

Keywords: Vaccine, Vaccination, Influenza, Flu, Externality, Health

*This paper has benefitted from suggestions by seminar participants including at the NBER Summer Institute (Health Economics), the SOLE Annual Meetings, the All-CA Labor Economics Conference, the ASHEcon Annual Conference, the IHEA World Congress, the WEAI Annual Conference, and various universities. This paper has also benefitted greatly from comments and suggestions from Desislava Byanova, David Chan, Olivier Deschênes, Mark Duggan, Melanie Guldi, Michelle Marcus, Sarah Reber, Peter Kuhn, Maya Rossin-Slater, Heather Royer, and Kosali Simon. Certain calculations use data from the National Center for Health Statistics multiple cause of death files (1994-2016), as compiled from data provided by the 57 vital statistics jurisdictions through the Vital Statistics Cooperative Program. All errors are my own.

1 Introduction

According to the Centers for Disease Control and Prevention (CDC), between 5% and 20% of the U.S. population are infected with influenza each year; these infections result in an average of approximately 200,000 hospitalizations and over 20,000 deaths.¹ Influenza is a vaccine-preventable disease, yet vaccination rates for influenza are substantially lower than vaccination rates for other vaccine-preventable diseases. This is largely due to the fact that the vaccine has to be received annually (and thus the cost of maintaining immunity is relatively high) and due to the lack of public policy incentivizing vaccination.

Vaccination serves as a canonical example of positive externalities in economics. Those who receive the vaccine incur some cost (monetary or otherwise) and experience a *private* benefit through the reduced risk of becoming ill; the *externality* benefit comes through the reduced risk of spreading the disease to others and the *social* benefit is the sum of the two. Because the benefits of vaccination are not fully internalized by the recipient, vaccines will be under-utilized relative to the social optimum in the absence of policy. This feature of vaccination has long been recognized by economists, and many theorists have considered how the socially optimal level of vaccination can be reached.² Achieving a social optimum requires information on both the marginal cost and the marginal social benefit of vaccination. While the private benefits of vaccination can be measured to some extent through the use of randomized controlled trials (RCTs), estimating the full extent of the social benefits requires an analysis at the population level.

This paper measures the marginal social benefit of influenza vaccination in two settings. First, I estimate the effects of state-level vaccination rates on influenza-related mortality and work absences in the United States. This portion of the analysis addresses the social benefits of vaccination in the general population. Second, I consider the efficiency gains to be had through targeted vaccination by examining a situation in which the externality benefits of vaccination are likely to be especially large. I analyze the impacts of county-level influenza vaccination mandates that apply to health care workers (HCWs) in California.

I measure the causal impacts of state-level vaccination rates by interacting pre-existing state-level differences in vaccination rates with year-to-year variation in the efficacy of the vaccine. Vaccine efficacy is measured as the extent to which the strains included in the season's vaccine match the strains that end up circulating. Mis-matches occur because of unpredictable genetic changes in the virus, and the prominence of these mis-matched viruses is not known until after vaccines have been distributed. Mis-matches provide an exogenous

¹Source: <http://www.cdc.gov/flu/about/qa/disease.htm>.

²For example: Stiglitz (1988); Brito et al. (1991); Francis (1997); Geoffard and Philipson (1997); Francis (2004); Boulier et al. (2007); Althouse et al. (2010); Manski (2010, 2017); Neilson and Xiao (2018).

source of variation in *effective* vaccination while allowing for *actual* vaccination rates to be held constant.

I find that higher vaccination rates lead to significant reductions in influenza-related mortality. Scaled nationally, I find that a one percentage point increase in the U.S. vaccination rate would result in approximately 807 fewer deaths per year in expectation. The mortality benefits primarily accrue to individuals 75 and older, but are mostly attributable to the vaccination of people under 75, suggesting substantial externalities. I also find that vaccination significantly reduces illness-related work absences. The estimates indicate that a one percentage point increase in the U.S. vaccination rate would result in approximately 15 million fewer work hours lost due to illness annually, in expectation. I find no impacts on either outcome during periods in which there is no influenza circulating and no impacts on outcomes that are implausibly related to influenza. In monetary terms, the estimates suggest that each vaccination confers at least \$64 in social benefits due to reduced mortality and \$98 in terms of reduced work absences.

The first component of the analysis is relevant to the vaccination of an average member of the population, yet there is likely substantial heterogeneity in social benefits of vaccination. In the second component of the analysis, I consider vaccination policy targeted at individuals with large potential externalities by exploiting the roll-out of county-level influenza vaccination mandates that apply to health care workers in California. Most of these mandates apply to all licensed health care facilities in a county, and thus there is potential for these mandates to reduce the spread of influenza both within the hospital (the unit of analysis) and in other health care settings (e.g., long-term care facilities). I find that these mandates increase hospital worker vaccination rates by 10.3 percentage points, reduce the number of influenza diagnoses for inpatient visits by 20.1%, and reduce the number of influenza diagnoses for outpatient emergency department visits by 8.1% during seasons with an effective vaccine. For inpatient visits, the impact is twice as large for influenza diagnoses that were not present at the time of admission (i.e., hospital-acquired infection). I estimate the marginal benefit of HCW vaccination in terms of health care cost savings to be \$131 per vaccination.

An exercise comparing the two components of the analysis suggests that health care worker vaccination is many times more effective at reducing the spread of influenza in comparison to vaccination in the general population. For both health care workers and the general population, the estimated marginal benefits of vaccination are large in comparison to the cost of vaccine administration, suggesting that programs that would increase vaccination at reasonable cost in either population are likely to be cost-effective. This is underscored by the fact that the analysis is limited to the outcomes for which data is available and as such the estimates do not reflect the benefits of vaccination along several other dimensions

including productivity losses for those who are sick at work (Pichler and Ziebarth, 2017) and long-term effects of in-utero exposure (Schwandt, 2017).

The primary contribution of this paper is to provide causal estimates of the social and externality benefits of influenza vaccination. While a large medical literature evaluates the benefits of influenza vaccination, much of the existing evidence on these benefits is derived from RCTs in which vaccination is randomized across individuals within a group, leaving no method for capturing externality effects.³ There are a limited number of studies in the RCT literature that directly evaluate externality effects by randomizing across groups rather than individuals (i.e., cluster RCTs). For example, Loeb et al. (2010) employ such a design, randomizing across isolated communities in Canada. In their study, influenza vaccinations were provided to children in the treatment communities and placebo vaccinations were provided to children in control communities. The authors find that vaccinating children led to reductions in laboratory-confirmed influenza for both children and adults in the treated communities, providing evidence of an externality benefit.

While it is possible to identify the presence of externalities in the context of an RCT, it is difficult to identify the effects of vaccination on severe and economically important outcomes such as mortality. The relative infrequency of the outcome would necessitate an extremely large-scale study; furthermore, ethical concerns over providing placebo vaccinations to high risk groups essentially relegates the study of *any* benefits (i.e., not only mortality) of influenza vaccination in the elderly population to an observational setting. The potential for bias in existing observational studies is large: a review of the evidence on vaccination in the elderly population noted implausibly large effects of vaccination on all-cause mortality, explaining that these results were likely due to, “systematic differences between the intervention and control arms” (Jefferson et al., 2010).

To my knowledge, there are few examples of papers that effectively circumvent this endogeneity issue; Ward (2014) is a notable exception. Ward (2014) uses exogenous variation in vaccine efficacy to evaluate the impacts of a regional influenza vaccination campaign in Ontario, Canada. The author finds that the program increased vaccination rates for non-elderly adults by approximately 10.8 percentage points (the post-treatment vaccination rate was approximately 33.3%) and resulted in a near elimination of influenza infection, a 92% reduction. The results suggest that Ontario reached a threshold level of vaccination beyond which the marginal benefits of vaccination fall to near zero. Models of influenza dynamics suggest the existence of such a threshold (Boulier et al., 2007), but the fact that an annual

³Reviews of this evidence are available from several sources, including the annual Recommendations of the Advisory Committee on Immunization Practices provided by the CDC (Grohskopf et al., 2014), a number of Cochrane reviews (Jefferson et al., 2010, 2012; Demicheli et al., 2014), and others (Osterholm et al., 2012).

epidemic is still experienced each year in the U.S. despite vaccination rates well above those during the study period in Ontario suggests that such a threshold has not been reached in the U.S. and that the results of the program in Ontario may have been specific to the location or period of analysis.

Similar to Ward (2014), my identification strategy relies on exogenous year-to-year variation in vaccine efficacy. My strategy, however, has the ability to exploit substantially more variation in vaccination rates and outcomes (across all U.S. states and 22 influenza seasons). As such, the average impacts that I estimate are not unduly influenced by the experience in any one region or time period. Compared to either Ward (2014) or a number of studies in the medical literature that do not tackle the endogeneity issues as directly (Jefferson et al., 2010), the estimates presented here are smaller in magnitude. That being said, the estimates are of plausible magnitude and still suggest that the social benefits of vaccination are substantial.

This paper also provides the first large-scale evidence on the impacts of influenza vaccination mandates for health care workers. This is an important contribution as such policies are actively being considered by regional public health departments. This is underscored by editorial articles published in several prominent medical journals that call for the adoption of such requirements (Stewart, 2009; Caplan, 2011; Hooper et al., 2014). The existing evidence on the benefits of such policies is derived from a small number of studies that assess the impacts of vaccination requirements primarily for employees of long-term care facilities. De Serres et al. (2017) critiques the findings of four recent studies by noting the implausibility of estimates based on non-specific outcomes. Another meta-analysis rates the overall quality of evidence on the subject as either “low” or “very low” (Thomas et al., 2016). My study has the advantages of using a highly specific outcome (hospital diagnoses for influenza), extremely large scale, and the ability to identify impacts in settings other than long-term care facilities.

While the specific contributions of this paper are described above, this paper also contributes to an empirical literature within economics that seeks to identify the economic impacts of influenza infection more generally. Much of this literature has focused on the effects of in-utero exposure to influenza on human capital development. Within this literature, most research has examined *pandemic* influenza (Almond and Mazumder, 2005; Almond, 2006; Kelly, 2011; Karlsson et al., 2014; Lin and Liu, 2014; Brown and Thomas, 2016), though exposure to *seasonal* influenza (the focus of this study) has been found to negatively impact health at birth and later-life outcomes as well (Currie and Schwandt, 2013; Schwandt, 2017). There are fewer studies that focus on the more contemporaneous impacts of influenza in the adult population, with notable exceptions studying vaccination (Ward, 2014) and other

factors that can influence the spread of influenza (Adda, 2016; Stoecker et al., 2016). This paper also contributes to a recent set of empirically-focused papers within economics that study various questions related to vaccination (Carpenter and Lawler, 2017; Lawler, 2017; Oster, 2018).

Finally, this paper contributes more generally to a literature in economics that seeks to empirically identify externality impacts in a variety of settings.⁴ Notably, there are few papers (exceptions above) that empirically identify externality or social impacts of vaccination, despite the fact that vaccines are often regarded as the “textbook” example of a positive externality (Stiglitz, 1988).

The remainder of the paper is structured as follows. Section 2 provides background information and a conceptual discussion that is helpful for interpreting the empirical analysis to follow. Section 3 (Part I) describes the analysis of vaccination rates in the general population, and Section 4 (Part II) describes the analysis of health care worker mandates in California. Section 5 offers a discussion and concludes.

2 Background

In this section, I provide a brief overview of several points regarding influenza and influenza vaccination that are necessary for interpreting the results of the empirical analysis. I also provide a conceptual discussion of the benefits of vaccination, focusing on the theoretical shape of the marginal benefit curves in the specific case of influenza vaccination.

2.1 Influenza and Influenza Vaccination

There are three key points regarding influenza for which I provide an overview in this section. First, I discuss the burden of influenza; specifically, it is important to understand the ways in which different groups are affected by the disease. Second, I discuss influenza vaccination, summarizing the current state of knowledge regarding vaccine efficacy. Third, I discuss in more detail the importance of vaccine match, as an understanding of the causes and consequences of vaccine mis-match are key to understanding the identification strategy used in the analysis to follow.

The total burden of influenza illness is large and crosses all demographic groups, though

⁴This literature is especially prominent in environmental economics, where many papers have sought to measure the impacts of pollution on a variety of outcomes (see Graff Zivin and Neidell (2013) for a review). Other examples include the evaluation of externality impacts of de-worming programs on health and schooling outcomes (Miguel and Kremer, 2004) and the estimation of displacement effects in job placement programs (Crépon et al., 2013).

there is substantial heterogeneity in how groups are affected. Infection is particularly severe in groups that the CDC identifies as at high risk for complication. These groups include: children under 5, adults 65 and older, pregnant women, and individuals with a range of chronic medical conditions. I focus on age as the dimension of heterogeneity to be studied in this paper. This discussion reflects the findings of the CDC’s Recommendation of the Advisory Committee on Immunization Practices, which summarizes the general findings from an extensive list of references (Grohskopf et al., 2014).

For children, influenza is responsible for large number of outpatient visits and hospitalizations, and this is especially true for infants.⁵ While outpatient visits and hospitalization are fairly common, death attributable to influenza among children is relatively rare. For non-elderly adults, influenza infection is typically less severe and less likely to result in hospitalization or death. While severe outcomes are less likely, the burden of influenza is still significant, often resulting in outpatient visits and worker absenteeism (Molinari et al., 2007). The majority of deaths related to influenza occur in individuals at least 65 years old. The CDC estimates that in 1976-2007, average annual deaths attributable to influenza were 21,098 for individuals 65 and older, 2,385 for individuals 19-64, and 124 for individuals under 19 (Thompson et al., 2010). As such, these estimates indicate that the 65 and older population account for approximately 90% of all influenza-related deaths. Due to difficulties in reporting and diagnosis, there is no consensus on the number of deaths that are caused by influenza in each year. Dushoff et al. (2006), for example, estimate a higher number of deaths compared to Thompson et al. (2010): an annual average of 41,400 for the period 1979-2001.

Influenza vaccine efficacy – the extent to which vaccination protects against laboratory-confirmed influenza – is determined by several factors. Vaccine match is an especially important factor, but it is important to note that even when the vaccine is perfectly matched it is not 100% effective. Vaccine efficacy also varies with age; diminished immune response among the elderly means that they are less able to create the antibodies needed to gain immunity. Estimates of vaccine efficacy in the prime-age population vary, though several studies find values in the range of 50-60% in a well-matched season (Demicheli et al., 2014; Grohskopf et al., 2014). Estimates of vaccine efficacy in the elderly population are more contentious, primarily due to the fact that ethical concerns over providing placebo vaccinations to high-risk populations limit the ability of researchers to use RCTs. There is some debate as to whether the vaccine provides any protective benefits among the elderly (Simonsen et

⁵Neuzil et al. (2000) find that influenza was responsible for an annual average of 6-15 outpatient visits per 100 children under 15. Additionally, Zhou et al. (2012) estimate annual influenza-related hospitalization rates (per 100,000) equal to 151 for infants, 38.8 for children aged 1-4, and 16.6 for individuals 5-49.

al., 2007), though a recent study reported by the CDC indicated efficacy of approximately 26% among people 65 and older (McLean et al., 2014) during a well-matched season.

Vaccine match – the degree to which the strains included in the vaccine match the strains that end up circulating – is an especially important determinant of vaccine efficacy. Understanding the identification strategy in the main analysis requires understanding the process by which a vaccine mis-match occurs. For the North American vaccine, this process begins in early Spring, when the World Health Organization convenes a meeting in order to make recommendations on the composition of the following season’s vaccine. The vaccine includes three (trivalent) or four (quadrivalent) strains, and the decision as to which strains to include in the vaccine is primarily based on which strains were circulating most recently.⁶ The Food and Drug Administration makes the ultimate decision regarding vaccine composition in the U.S., and vaccine composition is common across all states. Due to the time it takes to produce and distribute the vaccine, this decision must be made in early Spring so that vaccines can be administered in the Fall. The influenza virus itself undergoes constant genetic change (“antigenic drift”) such that there are always viruses in existence that are genetically distinct from the dominant strains; vaccines may not provide protection against these genetically distinct viruses. Significant mis-matches occur when one or more of these genetically distinct viruses becomes a dominant strain in a given season. Vaccine mis-matches are unpredictable prior to the start of influenza season. When a mis-match occurs, the non-matched virus strain is typically included in the following year’s vaccine. This means that matched and mis-matched strains do not represent a different set of virus strains, since mis-matched strains become matched strains in following seasons. It is also important to note that not only is the vaccine formulated well before influenza season begins, but individuals typically have no information on vaccine match at the time of vaccination. I provide direct evidence of this in Section 3.2.

Finally, a perfect match does not imply 100% efficacy, and a vaccine with zero matched strains does not imply 0% efficacy. Instead, the match rate serves as an exogenously determined proxy for vaccine efficacy. CDC estimates of vaccine efficacy across seasons provide insight into how efficacy varies with the match rate: over the period 2004/05 to 2015/16, estimates of efficacy range from 10% during a low match season to 60% during a high match season (Belongia et al., 2009; Treanor et al., 2012).

⁶The quadrivalent vaccine was introduced in 2012.

2.2 Marginal Social Benefits of Vaccination

Before moving on to the estimation, it is useful to consider a simple economic framework of externalities in the specific case of influenza. In this framework, there is a marginal private benefit of vaccination (MPB) and a marginal social benefit of vaccination (MSB). The MSB is assumed to be at least as large as the MPB at all points (i.e., the externality is non-negative). In a competitive equilibrium, consumers purchase vaccines such that the MPB equals the marginal private cost (MPC), and the vaccine is under-provided relative to a social optimum. The economic intuition is straightforward and is the basis for the analysis conducted in this paper. Considering the shape of the benefit curves in the specific context of influenza provides additional insight.

Boulier et al. (2007) combine basic externality theory with a workhorse model of disease dynamics (the susceptible-infected-removed “SIR” model) and parameterize the model to the case of influenza in order to derive theoretical predictions for the shape of the marginal benefit curves. Figure 1 produces a version of their result, allowing the MPB and MSB to depend on vaccine efficacy. I have plotted the MSB and MPB assuming 100% efficacy, as well as the MSB for 60% and 10% efficacy (efficacy is denoted E). The 60% and 10% curves are intended to reflect a well-matched and a poorly-matched vaccine, respectively.

Consider first the case of a perfectly effective vaccine (which is unrealistic for influenza, but instructive). The y-axis measures the number of infections such that at a vaccination rate of zero, the model predicts that an additional vaccination will prevent more than 1.5 infections in expectation; 0.5 infections are prevented in private benefits and the remainder are prevented in external benefits (the gap between the MPB and MSB represents the marginal externality). Measuring infections is equivalent to measuring the cost of disease if it assumed that the cost of infection is homogeneous and equal to one. The MPB is monotonically decreasing in the vaccination rate; this reflects the idea that an unvaccinated individual becomes less likely to be infected as the number of infections in the population decreases. While the MPB decreases, the MSB stays relatively flat until a threshold is reached. This threshold represents the point at which essentially all cases of influenza have been prevented and a seasonal epidemic fails to emerge (“herd immunity”). The shape of these curves prior to the threshold imply that neither the externality nor the social benefit of vaccination decreases prior to this point. Furthermore, it can be inferred that the U.S. is not beyond the threshold, as an influenza epidemic *does* emerge in each season. Current vaccination rates (approximately 43% in 2014) in combination with the persistence of an annual epidemic is at odds with the model that assumes $E = 100\%$ and predicts a threshold level of vaccination between 30% and 40%. At a more realistic $E = 60\%$, the MSB falls and the threshold increases. At $E = 10\%$ (i.e., a poor match season), the MSB is even smaller

and the threshold is never reached.

I caution that this model depends on a number of parameter choices that are difficult to estimate accurately, but considering the general shape implied by the model helps to guide the interpretation of the results to follow.⁷ Importantly, the model predicts that the externality makes up the majority of the social benefits of vaccination. The model also predicts relatively constant marginal social benefits of vaccination below the threshold, implying that estimates of the social benefits are unlikely to depend strongly on the level of vaccination. In other words, we should expect that the relationship between vaccination rates and the outcome is roughly linear until nearly all influenza cases are prevented.

Discussion of this framework also presents the opportunity to discuss potential heterogeneous impacts of vaccination (though heterogeneity is not explicitly built into the model). It is worth considering how two groups in particular may differ from the remainder of the population: the elderly population and health care workers. For the elderly, the cost of infection is high and vaccine efficacy is relatively low. These factors combined imply that the elderly benefit substantially from the vaccination of others and that a particularly large portion of the benefits to the elderly will operate through an externality. Health care workers (HCWs) are particularly interesting for two reasons. First, HCWs come in relatively frequent contact with infected individuals and thus vaccination is more likely to prevent infection in this group. Second, HCWs come in relatively frequent contact with individuals who have a high cost of infection (e.g., individuals with a compromised immune system), and thus the vaccination of HCWs may reduce the spread of infection precisely to those who would suffer the most severe consequences. These potential heterogeneous impacts motivate the focus on these groups in the empirical analyses to follow.

3 Part I: Vaccination in the General Population

3.1 Data

This analysis requires data on mortality by cause of death, illness-related work absences, influenza vaccination rates, the timing and magnitude of influenza activity, and the vaccine match rate. The unit of analysis is the state-year-month and the data cover the years 1994-2016.⁸ Because the analysis centers around influenza seasons, years are redefined as “flu-

⁷Important parameters include vaccine efficacy and the “contact number”, which is the number of additional infections that result from a single infection when the entire population is susceptible.

⁸The sample includes all states in the contiguous U.S.; Alaska, Hawaii and Washington D.C. are omitted in part because of small sample sizes in the data on vaccination rates, and in part because the extreme latitudes for Alaska and Hawaii relative to the rest of the U.S. affect the timing and magnitude of influenza season in these locations.

years”, running from July through June so that each flu-year represents a distinct influenza season. I often use “flu-years” and “influenza seasons” interchangeably. The data coverage is ultimately July 1994 through June 2016. In most specifications, the two flu-years affected by the 2009 H1N1 pandemic are omitted (2008/09 and 2009/10) so that the estimates represent seasonal influenza. Summary statistics for all data described here are provided in Table 1.

3.1.1 Outcomes: Mortality and Work Absences

Mortality data are derived from the multiple cause of death files from the National Vital Statistics System (NVSS). This is the restricted version of this data that includes state identifiers beyond 2005. It is important to note the use of multiple causes of death in classifying mortality as influenza-related. Dushoff et al. (2006) find that a large number of influenza-related deaths are excluded when only the underlying cause of death is used. Accordingly, deaths are classified by diagnosis if any of the (up to 21) diagnosis codes fall into the relevant category. Even using multiple causes of death, it is very rare for a death to be classified as specifically due to influenza. As such, the category with the highest level of specificity used in the analysis of mortality is deaths with any diagnosis for pneumonia/influenza (PI).⁹ Estimates are also presented that use only the primary cause of death instead of multiple causes for each of 34 mutually exclusive cause-of-death categories.

Data on illness absences are derived from the Current Population Survey (CPS) basic monthly files. Similar to Stearns and White (2018), the measure of illness absences is constructed using two questions posed to all individuals who report being employed. First, individuals who report being employed but absent from work for the entire reference week (i.e., worked zero hours) are asked the main reason for their absence. Second, individuals who are employed and at work during the reference week report both their usual hours worked and the number of hours actually worked in the reference week. Those who work less than 35 hours during the reference week but report that they usually work at least 35 hours per week are asked the main reason for working less than usual. Each of these two questions lists “own illness” as one possible reason for missing work and is the reason given for approximately 19% of absences (for both entire-week and partial-week absences). The main outcome of interest is the proportion of workers reporting an illness-related absence. In

⁹Because deaths due to influenza often occur as a result of complications or the exacerbation of pre-existing conditions, even PI deaths may exclude deaths that occurred as a result of influenza infection. As such, I also analyze deaths in two higher levels of aggregation: deaths with any respiratory or circulatory diagnosis, and all-cause deaths. Because it is highly unlikely that deaths without a respiratory or circulatory diagnosis occurred as a result of influenza infection, these non-respiratory/circulatory deaths are used as a falsification test. The ICD9 and ICD10 codes used to classify these diagnoses are as follows: Influenza (ICD9: 487-488, ICD10: J9-J11), Influenza/Pneumonia (ICD9: 480-488, ICD10: J9-J18), Respiratory/Circulatory (ICD9: 390-519, ICD10: I00-I99, J00-J99).

addition to illness-related absences, absences for other reasons are analyzed as falsification tests. Furthermore, for each worker reporting an absence, the survey also asks the number of hours missed; as an alternate measure, the average number of hours missed (for either illness or other reasons) is analyzed as well.

All measures of absence can be constructed only for individuals who work at least 35 hours per week and thus represent only full-time workers. A standard set of individual covariates are regressed out at the individual level prior to collapsing the residuals to the state-year-month level; covariates are indicators for gender, age (<20, 20-30, 30-40, 40-50, 50-60, >60), marital status (married, widowed/divorced/separated, never married), and education (less than high school, high school diploma, some college, college graduate).

3.1.2 Vaccination Rates

Data on state-level vaccination rates are obtained through the Behavioral Risk Factor Surveillance System (BRFSS). The BRFSS is a large-scale telephone survey that has been conducted in all states since 1993. The BRFSS asks whether each participant has received an influenza vaccination within the past 12 months.¹⁰ Due to the phrasing of the question, the season to which the vaccine applies is ambiguous for the months in which vaccines are distributed (primarily Sept.-Dec.). For example, in the month of October when many vaccines are received, an affirmative response may refer to a vaccine received in the current month for the upcoming influenza season, or to a vaccine received in the prior November or December for the prior season. For the main specification, I use all of the data and assume that responses in months Sept.-Dec. refer to the upcoming influenza season and responses in all other months (Jan.-Aug.) refer to the current or prior season. In an alternative specification, responses from months Sept.-Dec. are omitted to ensure that the season to which the response applies is not ambiguous.¹¹ All regressions are weighted by the number of BRFSS observations used to calculate the vaccination rate for the corresponding observation (unweighted regressions are presented in a robustness check).

¹⁰The exact phrasing of this question varies slightly from year to year. In more recent years, for instance, the survey asks about various types of vaccination (i.e., injection or spray). I classify each individual as having received an influenza vaccination if they received at least one dose of any type of influenza vaccine.

¹¹Information on vaccination was not collected for all states in survey years 1994, 1996, 1998 or 2000. In the main specification, vaccination rates for all influenza seasons are still calculated for all states; for example, survey responses from Sept.-Dec. 1999 are used to construct vaccination rates for the 99/00 season, and responses from Jan.-Aug. 2001 are used to calculate vaccination rates for the 00/01 season. In the alternative specification, only responses from Jan.-Aug. are used and vaccination rates cannot be calculated in all seasons for all states; as such, the following seasons are omitted from these regressions: 93/94, 95/96, 97/98, 99/00. As discussed in Section 3.3.4, the results are not qualitatively different using the alternate measure and excluding these seasons.

3.1.3 Match Rates

Data on the vaccine match are derived from annual influenza season summaries, which consist of data compiled from the CDC’s virologic surveillance system.¹² This system consists of laboratories located throughout the country that test respiratory specimens for the presence of any influenza virus and characterize viruses according to the exact strain. The data contain information on the number of viruses by strain and information indicating which strains the season’s vaccine protects against. The match rate for each season is defined as the percentage of characterized viruses that match the strains contained in that season’s vaccine. It is possible that the vaccine can offer some level of protection against strains that are not perfectly matched if the strain in the vaccine and the strain circulating are similar, and this information is indicated in the data. I construct two versions of the match rate, one in which strains are characterized as matched only if it is the exact strain contained in the vaccine, and one in which strains are characterized as matched if the vaccine offers some level of protection. The main specification uses the average of these two measures.¹³

3.1.4 Influenza Activity & Other Data

Data on the timing and magnitude of influenza activity are also obtained from the CDC’s virologic surveillance system. The primary measure of influenza activity is the percentage of tests that are positive for any type of influenza. Population data is required to construct mortality rates. Population by state, year and age are derived from the U.S. Census Bureau. Controls for temperature, humidity, and precipitation are included and derived from the Global Summary of the Day files.

3.2 Empirical Framework

Estimating the impacts of population-level vaccination rates is an empirically difficult task. To illustrate this difficulty, consider the following empirical equation:

$$Y_{sm_y} = \beta V_{s_y} + \Psi X_{sm_y} + \delta_{m_y} + \varepsilon_{sm_y} \quad (1)$$

¹²These data are available at: <http://www.cdc.gov/flu/weekly/pastreports.htm>.

¹³The exact process by which I calculate the match rate is slightly more complicated than I have laid out above. Each positive test received by the CDC is classified as either influenza A or influenza B. A subset of influenza A viruses are sub-typed (H1N1 or H2N3) and a subset of each subtype are then characterized to determine the exact strain. A subset of influenza B viruses are characterized to determine the exact strain. The annual summaries contain information on the number of total tests, the number of positive tests, the number of A and B viruses, the number of viruses sub-typed, the number of each subtype characterized, and the number of viruses belonging to a specific strain. Though relatively straightforward, I have developed a calculator that takes these numbers as inputs and outputs the match rate for each season. This calculator (along with all input data) is available upon request.

Y_{smy} measures the outcome in state s , flu-year y , and month m . The outcomes are either the PI mortality rate (per 100,000) or the average number of work hours lost due to illness. V_{sy} is the vaccination rate for state s and flu-year y , and δ_{my} are month-year fixed effects. X_{smy} represents a set of state-level time varying controls. In the main specification, this includes demographic controls and flexible controls for temperature, humidity and precipitation.¹⁴ The benefit of relating vaccination rates to outcomes at the state-level is that all within-state externalities are captured. The issue of course is that regressor in this equation (V_{sy}) is endogenous. It is certainly plausible, for example, that individuals in states that are more affected by influenza are more likely to get vaccinated. Indeed the estimate of β from the naive regression described above is 0.086 (s.e. = 0.024); taken causally, this would imply that higher vaccination rates lead to more influenza-related mortality. Controlling for state-specific factors may ameliorate this endogeneity issue: the inclusion of state-by-month fixed effects in the above equation produces a coefficient estimate of -0.012 (s.e. = 0.007). That being said, it is reasonable to argue that if fixed differences in vaccination rates across states are endogenous, then changes in vaccination rates are likely to be endogenous as well. This would be the case if, for example, states with rapidly increasing vaccination rates are more (or less) likely to pass effective health policy, or have individuals that are more (or less) likely to adopt other beneficial health behaviors. Such possibilities are far from implausible, and as such the strategy described below does not rely on potentially endogenous fixed differences or changes in vaccination rates for its identifying variation.

The strategy I use is to interact differences in vaccination rates with year-to-year variation in vaccine efficacy to generate plausibly exogenous variation in *effective* vaccination, while controlling for the actual vaccination rate. This strategy is illustrated in Figure 2. The top panel plots vaccination rates for two groups of states over time: a group of low-vaccination states (the five states with the lowest mean vaccination rate over time) and a group of high-vaccination states (the five states with the highest mean vaccination rate). The figure also plots the match rate in each influenza season. Note that vaccination rates in both groups evolve smoothly over time, and that the gap in vaccination rates evolves smoothly as well. Further note that there is no visual evidence indicating that vaccination rates are systematically different during high or low match seasons. The bottom panel plots the *effective* vaccination rate over time (i.e., the product of the actual vaccination rate and the match rate). During seasons in which the match is close to one, the gap in effective

¹⁴Demographics are population shares for 5-year age groups (0-4 to 75+), which vary at the state-by-year level. Weather controls vary at the state-year-month level. Temperature is expressed as the number of days in one of seven 10-degree mean temperature bins from <20F to >80F (60-70 omitted). Humidity is expressed as the number of days in three 5g/kg specific humidity bins from <5 to >20 (<5 omitted). Precipitation is expressed as a cubic in total monthly precipitation.

vaccination between low- and high-vaccination states is preserved. During seasons in which the match is poor, however, there is little difference between these states in effective vaccination. The identification strategy compares the difference in the outcome between low- and high-vaccination states in flu-years with a highly effective vaccine against the same difference in flu-years with a relatively ineffective vaccine. Equation (2) describes this difference-in-differences approach:

$$Y_{sm_y} = \gamma_1(V_{sy} * M_y) + \gamma_2 V_{sy} + \Psi X_{sm_y} + \delta_{sm} + \delta_{m_y} + \varepsilon_{sm_y} \quad (2)$$

In this equation, M_y is the match-rate, measured nationally for each flu-year. A match rate of zero implies that the vaccine is minimally effective whereas a match rate of one implies maximum efficacy. γ_2 absorbs the potentially endogenous component of the relationship between vaccination rates and the outcome by measuring the relationship between vaccination rates and the outcome in seasons in which the vaccine is minimally effective (i.e., zero match rate). γ_1 is the object of interest, and represents the differential effect of an increased vaccination rate between flu-years when the vaccine is at maximum versus minimum efficacy. Intuitively, γ_1 picks up the impact of *effective* vaccination (i.e., the causal effect of vaccination), but not the component of the relationship between vaccination rates and the outcome that persists in seasons when the vaccine is ineffective.

Because this strategy exploits variation in the match rate rather than variation in vaccination rates themselves, it is useful to consider the conditions under which Equation (2) identifies the effect of increased vaccination. $V_{sy} * M_y$ only identifies the true effect of vaccination under the condition that the vaccine is completely ineffective when the match rate is zero. This condition is not satisfied in reality, however, since the vaccine can provide some level of protection against non-matched strains. As such, $V_{sy} * M_y$ identifies a lower bound on the effect of vaccination, a concept which is illustrated in Figure A1 (and discussed in figure notes) using the model of infectious disease from Section 2. We proceed with the understanding that this strategy identifies a lower bound on the benefits of vaccination.

Identification relies on the assumption that the match rate is exogenous from year to year and that unobserved factors that are correlated with vaccination rates are unrelated to the match rate. One potential concern is whether it is possible for individuals to respond to the match rate in terms of vaccination behavior (e.g., choose not to receive a vaccine if the match is poor). Such behavior would introduce bias if there are differential responses across states with different vaccination rates. The process by which the strains are chosen for inclusion in the following season’s vaccine formulation, described in Section 2.1, supports the notion that the match rate is effectively random from year to year. Also supporting this notion is the fact that I find no evidence of serial correlation in the match rate, and no evidence of a

trend in the match rate over time.¹⁵ Furthermore, evidence described below suggests that there is limited scope for individuals to respond to the vaccine match.

From 2007 onward, the BRFSS has asked respondents not only whether they received a vaccination, but the month in which they were vaccinated. Figure A2 plots the average cumulative vaccination rate by month for the flu-years in which this data is available. Additionally, this figure displays average influenza intensity by month. Together, this figure shows that in a typical influenza season, nearly all vaccinations are administered before the onset of the season’s influenza outbreak. Because information on the vaccine’s match cannot be determined until a significant number of individuals are infected, this plot suggests that there is limited scope for responding to the match rate at all, much less differentially across states. That being said, while there is little scope for responding to match rates, it is not impossible.

Table 2 provides more direct evidence on the question of whether individuals are responding to match rates in terms of vaccination behavior. Column 1 reports estimates from a regression of the vaccination rate in a given state and influenza season on the match rate. The estimate is small and statistically indistinguishable from zero. The estimate implies that vaccination rates are 0.13 percentage points higher in 100% match seasons compared to 0% match seasons. The 95% confidence interval rules out a positive effect larger than 0.45 percentage points. In a more direct test of the identifying assumption, Column 2 presents estimates of a test for differential responses to vaccine match by interacting the match rate with the mean vaccination rate (over time) for each state; the results indicate no evidence of a differential response among states that tend to have higher or lower vaccination rates. To the extent that any concerns remain over the possibility of responses to the match rate in terms of vaccination behavior, additional specifications are estimated wherein only *pre-existing* variation in vaccination rates is employed. More specifically, the vaccination rate and each of its interactions are instrumented using vaccination rates from prior seasons. The estimates are insensitive to the choice of prior season, including the average vaccination rate over the prior three seasons, or a time-invariant vaccination rate defined as the average vaccination rate over the entire sample.

Further supporting a causal interpretation of the estimates is their robustness to a variety of specifications and falsification tests, described in detail in Section 3.3. One falsification test uses the idea that influenza vaccination will only have a causal effect on mortality and work absences during periods in which influenza is circulating. This idea is explicitly

¹⁵Both serial correlation and trends are tested at the flu-year level (the level at which the match rate is defined). A regression of M_y on M_{y-1} yields a coefficient estimate of -0.051 (s.e. = 0.263), and a regression of M_y on a time trend (flu-year) yields a coefficient estimate of 0.004 (s.e. = 0.010).

built into the main estimation strategy (described below), in which within-year variation in the timing and magnitude of influenza activity is exploited as a third source of variation. This can be done either by estimating Equation (2) separately for periods of high and low influenza activity, or more formally in a triple-difference approach. While both versions are presented, the advantage of the triple-difference approach is that it allows for the use of a continuous measure of influenza activity, which more precisely focuses on the periods in which the largest impacts are expected.

$$\begin{aligned}
Y_{sm_y} = & \phi_1(V_{sy} * M_y * A_m) + \phi_2(V_{sy} * M_y) \\
& + \phi_3(V_{sy} * A_m) + \phi_4 V_{sy} + \Psi X_{sm_y} + \delta_{sm} + \delta_{m_y} + \varepsilon_{sm_y}
\end{aligned}
\tag{3}$$

Influenza activity, A_m , can be defined in different ways. One option is to define A_m as an indicator representing influenza season months (i.e., December-March). Another approach is to more precisely focus on periods in which either the largest or smallest impacts are expected, and define A_m as average monthly influenza activity. Although these two methods obtain similar results in terms of magnitude, the latter is preferred for precision. The influenza activity measure is an index that is scaled to equal one during the month with maximum average influenza activity. A_m varies only across calendar months; within a month, there is no variation across years or states. The measure takes its largest value (equal to one by construction) during February and its smallest value (0.052) during July. Note that the main effect for match (M_y), the main effect for activity (A_m), and their interaction ($M_y * A_m$) are implicitly included in the month-by-year fixed effects. In Equation (3), ϕ_1 measures the difference in the causal effect of vaccination (γ_1 from Equation (2)) between periods of maximum expected influenza activity and periods of zero expected activity. The coefficient ϕ_2 represents a falsification test: the causal effect of vaccination during periods in which essentially zero influenza is circulating.

Equation (3) includes both month-by-year fixed effects and state-by-month fixed effects. The state-by-month fixed effects not only account for any fixed state-specific factors, but allow seasonality in the outcome to vary by state. The identification strategy does not rely on the inclusion of state fixed effects, but inclusion of state-by-month fixed effects is preferred for precision. The main specifications include controls for 5-year age shares and flexible controls for temperature, precipitation, and humidity (Barreca and Shimshack, 2012). Rather than including a more exhaustive list of state-by-year controls, I recognize that Equation (3) allows for the inclusion of state-by-year fixed effects and demonstrate that the results are not sensitive to their inclusion. These results are presented alongside a range

of other specifications in Table A1. In all specifications, standard errors are clustered at the state level.

3.3 Results

In interpreting the estimates, note that the three regressors (*Vaccination*, *Match*, and *Activity*) are all continuous measures. *Vaccination* is measured on a scale of zero to 100 such that the estimates can be interpreted as a one percentage point increase in the vaccination rate, while the other two regressors are measured on a scale of zero to one.

3.3.1 Main Results

Table 3 provides estimates of the difference-in-differences equation described in Equation (2). These estimates are presented for both the all-age PI mortality rate per 100,000 population (Panel A) and the proportion of workers absent due to illness (Panel B). Column 1 includes all months, and Columns 2-3 compare influenza season months (December-March) to non-season months (April-November).

First, let us consider the coefficient estimates for *Vacc*, which represent the relationship between vaccination rates and the outcomes when the match rate is zero. These estimates are positive and larger during influenza season months. This suggests that states with higher influenza-related mortality or more influenza-related work absences tend to have higher vaccination rates. In other words, these estimates pick up the potentially endogenous component of the relationship between vaccination and the outcomes.

The coefficient estimates for the interaction term are the objects of interest. These estimates represent the differential effect of higher vaccination rates between flu-years in which the match is equal to one relative to flu-years in which the match is equal to zero. For PI mortality, the all-month estimate in Column 1 is negative and significant; the point estimate of -0.043 implies that a one percentage point increase in the vaccination rate decreases PI mortality by 0.043 per 100,000 population in flu-years when the vaccine is perfectly matched relative to flu-years when the match is zero. Magnitudes will be discussed in greater detail in the discussion of Table 4 below. For work absences, the all-month estimate in Column 1 is also negative, but it is not statistically different from zero. The point estimate of -0.00010 implies that a one percentage point increase in the vaccination rate decreases the proportion of workers absent due to illness by 0.00010 percentage points (an approximate 0.4% decrease). The lack of precision and relatively small magnitude of these estimates in part motivates analyzing months in which the benefits are expected to be largest.

For both PI mortality and work absences, the general patterns in comparing influenza

season months to other months are similar: the estimates are larger in magnitude and statistically significant during influenza season and smaller in magnitude during other months. The estimate for PI mortality during non influenza season months is negative and significant, though this is not necessarily unexpected since the estimates represent months with relatively little influenza activity rather than zero activity.

The triple-difference model formalizes the comparison of months with high and low expected influenza activity. This formalization has the advantages of (1) utilizing the entire distribution of influenza activity rather than a binary classification of months in to “High” and “Low” activity, and (2) providing a falsification check in the form of a coefficient estimate that represents the causal effect of vaccination during periods of *zero* measured influenza activity. The triple difference estimates for both PI mortality and work absences are reported in Table 4. Column 2 in this table represents the main specification, but the table presents two additional specifications as well: Column 1 represents a more parsimonious specification that omits the state-by-month fixed effects and Column 3 builds on the main specification by instrumenting for the vaccination rate (and all interactions) using the average vaccination rate over three seasons prior (and all interactions).¹⁶

Next consider the magnitude of the estimates in Column 2 of Table 4. Note that all subsequent analyses and calculations are based off of this specification. For both PI mortality and illness absences, the coefficient on the $Vacc \times Match \times Activity$ interaction is negative and significant at the 5% level, while the coefficient on the $Vacc \times Match$ interaction is small and indistinguishable from zero. The implication is that influenza vaccination rates decrease both PI mortality and illness absences during periods of high influenza activity, but have no impact on the outcomes when measured influenza activity is zero. The point estimate of -0.073 on PI mortality implies that a one percentage point increase in the influenza vaccination rate will decrease the PI mortality rate by 0.073 per 100,000 individuals during months with maximum expected influenza activity relative to months with no expected influenza activity, and during seasons in which the vaccine is perfectly matched relative to seasons in which the vaccine is poorly matched. Similarly, the coefficient estimate of -0.00048 on illness absences implies that a one percentage point increase in the influenza vaccination rate will decrease the proportion of full-time workers absent due to illness by 0.00048 percentage points under the same conditions.

Because these interpretations are somewhat nonintuitive, I also report an estimate of the “Expected Annual Benefit” of vaccination. To calculate the expected annual benefit, I first

¹⁶Additional specifications that vary the fixed effects and use of the instrument are presented in Table A1, and alternative definitions of the instrument are explored in Table A2. For ease of interpretation, these alternative specifications are presented graphically alongside a variety of other robustness checks (discussed later). Across all alternative specifications the estimates are of comparable magnitude.

calculate expected monthly benefits and then sum across months (described in table notes). These estimates imply that in a population of 100,000 individuals, a one percentage point increase in the vaccination rate (i.e., 1,000 additional vaccines) would decrease PI mortality by 0.249 and decrease work hours lost to illness among full-time workers by 4,716 hours in expectation. Put differently, the estimates imply that 4,016 vaccinations are required to save one life and 1.70 vaccinations are required to save one 8-hour work day. Finally, suppose the estimates are scaled to the size of the U.S. population: a one percentage point increase in the vaccination rate for one year would result in 807 fewer deaths and 15 million fewer work hours lost in expectation. For mortality, this is a substantial, albeit plausible number given that estimates of average annual deaths due to influenza lie in the range of approximately 20,000-40,000.

Before moving on, it is useful to explore alternative definitions of the outcomes for the purpose of ensuring that the estimated benefits are concentrated in the outcomes that are expected to be most heavily influenced. Consider the impacts on mortality by 34 standard cause of death categories in Figure 3. The cause-of-death categorization here uses the underlying cause of death only so that each category is mutually exclusive (as opposed to using multiple causes of death to define PI deaths as in the main analysis). As such, it is possible that deaths in non-PI categories have a secondary diagnosis for PI and so categories other than PI should not be considered as true falsification tests. It is reassuring nonetheless that the largest and most highly significant coefficient estimate is for the PI category.¹⁷

3.3.2 Age Heterogeneity and Decomposing Externality Effects

To this point, all estimates have represented the social benefits of vaccination. We next consider evaluating mortality benefits by age with the ultimate goal of disentangling private and external benefits. Age-specific estimates for mortality are provided for five age groups (infants under 1, 1-9, 10-64, 65-74, and 75+) in Panel A of Table 5. Age-specific mortality rates are calculated using the total state population in the denominator rather than the age-specific population so that the estimates can be interpreted as an accounting of the total

¹⁷Table A3 presents alternate definitions of mortality based on multiple causes of death as well as alternative definitions for work absences. In both cases, the main estimates are provided for reference. For mortality, results are provided for two broader definitions of influenza-related mortality (respiratory/circulatory, and all-cause), and for a falsification test (non respiratory/circulatory mortality). The standard errors on these broad categories are too large to draw meaningful conclusions. Alternative definitions for work absences are provided as well. First is average hours absent due to illness rather than proportion absent; this estimate is negative, marginally significant, and of similar magnitude to the main estimate (when scaled by the number of hours per absence). Also provided are estimates for both average hours absent and the proportion absent for reasons other than illness as falsification checks; these estimates are small relative to the mean and insignificant.

benefits.¹⁸ The age-specific estimates indicate essentially all of the mortality benefits accrue to the 75+ population (in fact the coefficient estimate is almost identical to the all-age estimate). This is not a surprising result given that estimates of influenza-related mortality are heavily concentrated among the elderly population: Grohskopf et al. (2014) reports that the 65+ population accounted for over 90% of all influenza-related deaths between 1976 and 2007.

To evaluate the extent to which the mortality benefits of influenza vaccination operate through an externality effect, I use the fact that the vast majority of benefits accrue to individuals who are at least 75 years of age and separately estimate the effects of vaccination rates for individuals who are either within or outside of that age group.¹⁹ More specifically, the following equation is estimated:

$$Y_{smy}^{O75} = \psi(V_{sy}^{O75} * M_y * A_m) + \omega(V_{sy}^{U75} * M_y * A_m) \quad (4)$$

+ Other Interactions & Controls + ε_{smy}

In Equation (4), the full set of interactions described in Equation (3) for both people under 75 and people at least 75 are included. As such, ψ represents a combination of direct and externality effects, where the externality effects are limited to capturing the spread of influenza among people within the 75 and older group. The coefficient ω represents the effect of vaccination among people under 75 on influenza-related mortality for individuals who are at least 75; this represents a pure externality effect. The results of this exercise are presented in Panel B of Table 4. The results indicate that essentially all of the mortality benefits of influenza vaccination operate through an externality channel, though it cannot be rejected that the estimates for ψ and ω are equal. Nevertheless, this finding accords with theoretical predictions presented in Figure 1 which indicate the majority of the social benefits of vaccination operate through an externality. Furthermore, the theoretical prediction does not taken into account the relatively low efficacy of influenza vaccination in older individuals. Taking this heterogeneity into account, the model would predict an even greater proportion of the benefits accruing to the elderly population operate through an externality, consistent with the empirical result presented here.

¹⁸An additional set of estimates is provided in Table A4 that uses the age-specific population in constructing mortality rates with very similar results.

¹⁹Age-specific vaccination rates are constructed using the BRFSS, which is individual-level data with information on each respondent's age.

3.3.3 Monetizing Benefits

Any policy aimed at increasing influenza vaccination take-up should weigh the costs and benefits of doing so; the goal of this section is to provide monetary estimates of the marginal social benefits of vaccination in terms of both mortality and work absences.

The monetary benefits of influenza vaccination in terms of mortality depend on the value of a statistical life (VSL). Because the mortality benefits are concentrated among individuals at least 75 years of age, it is especially important that the VSL is age-adjusted. I use the method of Murphy and Topel (2006), who develop a framework for estimating the value of remaining life given a standard VSL figure that is evaluated using mortality risk reductions from working-age adults. I apply two such figures: estimates from Ashenfelter and Greenstone (2004) of \$2.3 million (denoted “AG”) as a lower bound, and the current EPA standard of \$8.8 million as an upper bound.²⁰ The Murphy and Topel (2006) framework provides VSL estimates for single years of age; I follow Deschênes et al. (2017) to calculate a VSL estimate for the 75+ age group, taking a weighted average of single-year VSL estimates where the weight is the share of deaths from each single-year of age.

Table 6 presents the expected annual benefits of a policy that increases the national influenza vaccination rate by one percentage point. The expected annual number of deaths avoided for the 75+ group (807.3 deaths) is multiplied by the age-adjusted VSL.²¹ Benefits per vaccination are determined by dividing the total benefit by 3.23 million (1% of the 2016 U.S. population), as that is the number of additional vaccinations required to achieve the corresponding total benefit. The benefits per vaccination are calculated to be \$64 using the AG VSL, and \$244 using the EPA VSL. Recall that the results from Table 5 suggest that these benefits operate primarily through an externality channel.

Focusing only on mortality reductions in evaluating monetary benefits of vaccination is problematic for at least two reasons. First, monetizing life is inherently controversial among the general public. Second, it is possible that some of the deaths avoided due to increased vaccination represent a population suffering from comorbidities; if this is the case then the VSL estimates (which are only adjusted for age and not comorbidities) would be overstated. As such, it is advantageous to also present benefits of vaccination that are subject to less controversy and require fewer assumptions in monetizing. The process of calculating monetary benefits for work absence reductions is somewhat more straightforward: the expected annual number of hours saved for the U.S. population is multiplied by the median hourly wage. The calculation suggests that each vaccine confers benefits equal to \$98 in terms of reduced work absences among full-time workers, in expectation.

²⁰Each VSL figure is reported in 2016\$.

²¹These calculations are similar to the “Expected Annual Benefits” calculations in Table 4.

3.3.4 Robustness Checks

This section provides a brief summary of a number of robustness checks. All results are presented in the Online Appendix, and specifications are discussed more fully in table and figure notes. Many of these robustness checks are summarized in Figure A3, which displays coefficient estimates and confidence intervals for a range of specifications.

Some of these robustness checks have already been mentioned. Table A1 presents a series of specifications demonstrating that the estimates are insensitive to the choice of fixed effects and the use of the instrumental variables strategy. Table A2 demonstrates that the estimates are insensitive to various definitions of the instrument, including defining the instrument as a time-invariant vaccination rate (the state-level average over the sample). Table A4 demonstrates that the age-specific estimates are not sensitive to defining mortality rates using the age-specific population.

Potential nonlinearities in the marginal effect of vaccination are explored in Figure A4. This exercise shows little evidence that the marginal benefit of vaccination changes substantially over the observed distribution of vaccination rates (consistent with the theoretical model in Figure 1). The possibility of lagged effects on mortality or work absences is examined in Table A5. The reported estimates represent the sum of contemporaneous and lagged impacts (for multiple lag lengths), and the stability of the estimates suggest that the vast majority of the impacts operate contemporaneously, though the standard errors grow considerably as more lags are added.

Table A6 presents several additional sensitivity checks, all of which are also summarized visually in Figure A3. Column 1 includes the two flu-years affected by the H1N1 pandemic. Column 2 uses an influenza season indicator in place of influenza activity in the triple difference specification; a “scaled effect” is provided for comparability with the main estimates. Column 3 uses a limited sample that ends in the 2007/08 flu-year to eliminate the influence of three factors: (1) the H1N1 pandemic, (2) the development of the high-dose vaccine in 2009, and (3) the development of the quadrivalent vaccine in 2012. The estimate for mortality is smaller in magnitude when seasons beyond 2007/08 are omitted (the estimate for work absences is slightly larger), though it should be cautioned that much of the identifying variation is thrown out when these seasons are omitted. Column 4 uses the sample 1998/99-2014/15 in which region-specific information on the match rate and influenza activity are available, and allows the match rate to be defined as region-specific. Column 5 provides unweighted estimates. Column 6 uses an alternative calculation for vaccination rates, described in Section 3.1.2. Column 7 uses a strict definition of vaccine match, described in Section 3.1.3.

4 Part II: Health Care Worker Mandates

The analysis conducted in Section 3 was intended to estimate the benefits of influenza vaccination in the general population. The estimates are relevant to a policy that would increase vaccination among those who are on the margin of the decision to receive a vaccination. In this section, I recognize that there is likely to be substantial heterogeneity in benefits depending on who receives the vaccine. Health care workers (HCWs) come in relatively frequent contact with infected individuals *and* individuals whose cost of infection is high. As such, HCWs are a group for whom the external benefits of vaccination are likely to be particularly large. I examine the effects of mandates requiring health care workers be vaccinated against influenza on the outcomes of hospital patients in counties and hospitals subject to the mandates. To begin, I describe the institutional background.

4.1 Institutional Background

On September 28, 2006, the Governor of California signed into law Senate Bill 739, requiring that health facilities implement various measures to protect against the spread of infection. This law required that all health facilities offer free vaccinations to employees and required that they sign a statement declaring that he or she had declined vaccination if that was the case. Though detailed data on vaccination rates prior to this policy are not available, it is likely that these policies increased vaccination rates of HCWs. This means that baseline levels of vaccination are relatively high by the time the first mandates go into effect in 2009. This law also required all hospitals to report to the California Department of Public Health (CDPH) on the percentage of HCWs vaccinated against influenza in each season, allowing for estimation of first-stage impacts of vaccination mandates on vaccination rates.

In May of 2009, the H1N1 pandemic began. In response to the pandemic, a small number of individual hospitals began requiring influenza vaccination for their workers. After the 2009 pandemic, these hospitals continued requiring annual influenza vaccinations for their workers and in following influenza seasons several other hospitals began introducing their own mandates. Beginning in the 2011/12 season, counties implemented county-wide vaccination mandates, and in each season since more counties have followed.²² In the 2015/16 season, 341 hospitals in California (over 75% of hospitals) were subject to a mandate. Of these 341 hospitals, 27 implemented their own mandates and the remainder were subject to mandates imposed at the county level. In other words, the vast majority (over 90%) of the hospitals

²²Implementation dates for all hospital- and county-level policies are described in Table B1.

ever subject to a mandate did not choose to impose it.²³

The county-level policies were not all implemented in exactly the same fashion. Specifically, a limited number of these policies applied only to hospitals, while most mandates applied much more broadly. Typically, all licensed health care facilities would be subject to these more broad mandates. Figure B1 maps the implementation of both hospital and county-level mandates over time and distinguishes between the type of county-level mandate. Because the vast majority of mandated hospitals are subject to county-level mandates that apply beyond the hospital, these are the policies on which I focus. The main outcomes are measured at the hospital level, but because these mandates apply more broadly, it is not necessarily the case that an observed infection was transmitted within the hospital. It is possible that these mandates affect the transmission of influenza in non-hospital health care settings and in the community at large if HCWs act as important vectors for disease.

4.2 Data

To estimate the effects of HCW vaccination mandates on patient outcomes, I make use of data on the timing of the mandates, vaccination rates for HCWs and hospital-level outcomes. In the main analysis, the data cover flu-years 2007/08 through 2015/16.

4.2.1 Mandates & Vaccination Rates

Data on the timing of mandates is compiled from several sources. Information on hospital-level mandates comes largely from the Immunization Action Coalition (IAC), a non-profit immunization activist group that lists health care organizations across the U.S. that mandate influenza vaccination and the dates of implementation.²⁴ CDPH maintains a list of county-level mandates with implementation dates, but the list is not completely accurate with respect to the implementation dates. Through a process of searching for county-level public health orders and identifying the initial date of implementation, I have either verified or amended the dates of nearly all counties on the list provided by CDPH. For an example county public health order (Alameda County in 2013), see Figure B2.²⁵ Summary statistics in Table 1 indicate the number of hospitals subject to mandates in each flu-year.

As required by California law, all licensed hospitals report information on the vaccination status of their workers to CDPH for each season. This information is compiled in their annual

²³To ensure results are not driven by hospitals that select into treatment or by the experience during the H1N1 pandemic, Table B2 demonstrates that the results are robust to (1) excluding the hospitals that adopt their own mandates, and (2) excluding flu-years affected by the pandemic.

²⁴Source:<http://www.immunize.org/honor-roll/influenza-mandates/>.

²⁵I have compiled a number of public health orders for other counties and years, and these documents are available upon request.

Hospital Employee Influenza Vaccination Reports. Though all hospitals provide information on vaccination rates, the within-hospital response rate is not 100%. Reporting in the first season in which reporting was required (2008/09) was particularly poor, and so data from this season is omitted. For the remaining flu-years, the main first-stage estimates use only hospitals that have response rates of at least 90% in all flu-years, though the estimates are not sensitive to this restriction.

4.2.2 Hospital Patient Outcomes

The primary data source on outcomes are two restricted data files on the universe of inpatient hospital admissions and outpatient ED visits in California between 2005 and 2016, obtained through California’s Office of Statewide Health Planning and Development (OSHPD). In the analysis to follow, inpatient admissions and outpatient ED visits are analyzed separately, though I refer to any visit to the hospital (inpatient admission or ED visit) simply as a “visit”. Note that ED visits often result in hospital admission; to avoid double-counting, ED visits are dropped if the patient is transferred to another health care facility (i.e., all ED visits are outpatient).

Unlike the mortality data in which influenza is rarely indicated as a cause of death, hospital patients routinely receive diagnoses specifically for influenza, allowing the outcome measure to be more specific. The primary outcomes of interest are the number of inpatient admissions and the number of outpatient ED visits with any diagnosis for influenza; admissions are classified as such if any of up to 25 diagnoses are for influenza.

The inpatient data include a number additional features that are utilized. One particularly useful feature is that each diagnosis includes an indicator for whether it was present at the time of admission, allowing me to focus specifically on hospital-acquired infection. Certain outcomes are less specific to influenza, but may be significantly affected during periods of very high influenza activity. These include average length of stay, average hospital charges, and the in-hospital death rate.

In addition to these hospital-level measures, I also examine PI mortality, which is observed at the county level using restricted data files from the NVSS. Summary statistics for all outcomes are presented in Table 1.

Finally, the analysis also uses data on county-level time varying covariates: uninsurance rates from Small Area Health Insurance Estimates (SAHIE), unemployment rates from Local Area Unemployment Statistics (LAU), per-capita income and government transfers from the Regional Economic Information System (REIS), and 5-year population age shares from the Surveillance, Epidemiology and End Results system (SEER).

4.3 Empirical Framework

I estimate the impacts of vaccination requirements using a standard difference-in-differences (DD) framework that exploits quasi-experimental variation in the timing of mandates. Similar to the general population analysis, I also use a triple-difference framework that additionally exploits the timing and magnitude of influenza activity for outcomes that are not specific to influenza. Because vaccination rates (the first stage) are measured annually, and because there is no variation in influenza-specific diagnoses (the main outcomes) during months with no influenza circulating, the triple-difference strategy is not appropriate for these outcomes. Consider the following DD equation to be estimated at the annual level:

$$Y_{hy} = \alpha + \pi \text{Required}_{hy} + \Psi X_{cy} + \delta_h + \delta_y + \varepsilon_{hy} \quad (5)$$

In Equation (5), Y_{hy} represents either vaccination rates (first stage), or the number of influenza diagnoses (reduced form) at hospital h in flu-year y . Required_{hy} is a variable indicating whether there is a vaccination requirement in effect. X_{cy} is a vector of county-level time-varying covariates, and δ_h and δ_y are hospital and flu-year fixed effects.²⁶ The coefficient of interest, π , is identified under the assumption that variation in the timing of the mandates is uncorrelated with other unobserved time-varying determinants of the outcomes.

While the identifying assumption is fundamentally un-testable, I provide evidence from indirect tests that support the assumption. Most importantly, in the discussion of results I provide an event study version of Equation (5). This exercise provides strong evidence that changes in the outcomes coincide precisely with the implementation of the policy, and that the treatment effects are not identified off of differential trends between treatment and control hospitals. Furthermore, the event study provides information on the necessity of county- or hospital-specific time trends, which are included in some specifications.

Because influenza diagnoses represent a highly specific outcome, there are many hospital-by-flu-year cells with the outcome equal to zero. Given the count nature of the data, and the over-dispersion indicated in Table 1 (i.e., the variance is greater than the mean), a negative

²⁶Covariates include the health un-insurance rate, the unemployment rate, per capita income, per capita government transfers, and 5-year population age shares (0-4 to 75+).

binomial model is used when influenza-specific diagnoses are the outcome.²⁷ With a count model it is important to allow the probability of an event to occur (i.e., an influenza-related diagnosis) to differ by hospital size, which varies considerably across the sample. This is done through the use of an exposure variable, which is set to be the mean annual number of all-cause visits in 2005/06-2006/07 (prior to the period of analysis). The first stage is estimated via OLS.

For outcomes that vary across all months of the year, the preferred specification is a triple-differences, estimated at the monthly level and taking the following form:

$$Y_{hmy} = \alpha + \theta_1(\text{Required}_{hmy} * \text{Activity}_{my}) + \theta_2\text{Required}_{hmy} + \Psi X_{cmy} + \delta_h + \delta_{my} + \varepsilon_{hmy} \quad (6)$$

In Equation (6), the policy indicator is interacted with an index of influenza activity, Activity_{my} . This measure of influenza activity is distinct from the measure used in Part I of the paper in that it measures actual influenza activity in a given month-year rather than the monthly average across years. In the analysis of health care workers, it is more reasonable to argue that influenza activity, measured nationally, is exogenously determined. The use of actual influenza activity allows the estimates to capture the effects of vaccination mandates at the exact time the largest impacts would be expected, maximizing the power of the estimates. The index ranges from zero to one, where one is the maximum observed value in the sample. The main effect for activity is absorbed by the month-year fixed effects. θ_1 measures the effect of vaccination mandates during a time of peak influenza activity relative to a period with zero influenza activity. θ_2 measures the effect of influenza vaccination mandates during times of very low influenza activity and is expected to be near zero. Because the outcomes

²⁷There are several possible count models available, and in the case of panel data requiring fixed effects (as here) the choice is not trivial (See Cameron and Trivedi (2013a,b) for a review of count models in general and specifically for panel data.). The workhorse count model that allows for fixed effects is the Poisson fixed effects estimator (Hausman et al., 1984; Wooldridge, 1999); this estimator, unlike many nonlinear models, provides consistent estimates of the slope parameters in the presence of fixed effects. A deficiency of the Poisson model, however, is that it assumes that the variance and mean of the outcome are equal (i.e., equi-dispersion). The usual solution is to use a negative binomial in place of a Poisson model, which allows for over-dispersion in the data. Hausman et al. (1984) offer a fixed-effects version of the negative binomial, but subsequent work has pointed out that this model requires an additional and often unrealistic assumption regarding the relationship between the fixed effects and the over-dispersion parameter (Allison and Waterman, 2002; Guimaraes, 2008). An alternative strategy is to estimate a standard negative binomial model with a full set of indicators as fixed effects. In nonlinear models using short panels, this leads to biased and inconsistent estimates of the slope parameters due to an incidental parameters problem. That being said, Allison and Waterman (2002) provide evidence from Monte Carlo simulations that suggests little bias resulting from the incidental parameters problem in the case of the negative binomial model with indicator as fixed effects. I adopt the negative binomial with indicators as fixed effects as the main specification, though the results are not sensitive to the choice of count model.

(average charges, average length of stay, and the in-hospital death rate) are not counts, Equation (6) is estimated via OLS.

The data on mortality are derived from a different data source than the hospital-level measures. The mortality data are only available at the county level and as such estimates are conducted at the county level (with county fixed effects), and the outcome is PI deaths per 100,000 population. In all models, standard errors are clustered at the county level.

4.4 Results

4.4.1 First Stage

The main result for the first stage is illustrated as an event study in Figure 4. This figure shows that there is little evidence of differential trends between hospitals that do and do not adopt vaccination mandates prior to implementation. In the first flu-year of implementation, vaccination rates increase sharply and remain relatively flat thereafter. The result is shown using a single treatment indicator in Table B3, revealing a highly significant coefficient estimate of 10.3 percentage points in the preferred specification (Column 2).²⁸

It is important to keep in mind that the first-stage estimates only represent vaccination rates for hospital workers. This is especially important in considering the county-level requirements, which apply far more broadly than to just hospital workers. Because vaccination rates for other HCWs are not observed, results are displayed in the remainder of the paper as reduced-form policy estimates rather than in an IV framework. That being said, there is reason to believe that the first-stage effect for non-hospital HCWs may be larger than that of hospital HCWs. The CDC conducts an online survey that provides national estimates of influenza vaccination rates for HCWs by place of work and whether or not vaccinations are required. The 2016/17 survey indicated vaccination rates of 82.6%, 68.7%, 58.5% and 56.2% for HCWs in hospital, ambulatory care/physician’s office, long-term care, and other settings, respectively. Because hospital workers tend to have the highest baseline vaccination rate, it is likely that influenza vaccination requirements have a larger effect on workers in settings with a lower baseline level. Indeed, the smallest gap in the vaccination rate between required and non-required settings is for hospital workers. Specifically, the required/non-required gap is equal to 19.3 percentage points (p.p.), 38.7 p.p., 39.0 p.p., and 37.4 p.p. for HCWs in hospital, ambulatory care/physician’s office, long-term care, and other settings, respectively. While this strongly suggests that effect of mandates on vaccination rates is larger in non-hospital settings, it is also possible that enforcement is weaker in non-hospital settings and

²⁸This table also shows that the estimate is robust to using less restrictive sample selection criteria on data quality, or even using all hospitals regardless of data quality.

this cannot be tested this with available data.

To get a rough estimate of the number of additional vaccinations received as a result of these mandates, consider a hypothetical state-wide mandate that affected all workers in licensed health care facilities in California. The Bureau of Labor Statistics indicates there were approximately 1.18 million workers in industries plausibly affected by these mandates in California during 2015.²⁹ Assuming the mandates affected all HCWs in the same way hospital workers were affected, multiplying by the first-stage estimate of 0.103 implies 121,154 additional vaccinations.

4.4.2 Influenza Diagnoses

Before discussing magnitudes, first consider a number of specifications for the reduced-form estimates presented in Table 7; Panel A reports inpatient admissions and Panel B reports outpatient ED visits. Column 1 represents estimates with no included trends, covariates, or sample restrictions. The estimates are in the expected negative direction for both outcomes, but only marginally significant for inpatient admissions and insignificant for ED visits. While there is not enough variation in vaccine efficacy over the sample period to incorporate this into the model as in Part I, there is one flu-year in the sample (2014/15) that had a poor vaccine match and thus an ineffective vaccine. Dropping this flu-year from the sample results in larger coefficient estimates, reported in Column 2. This flu-year is omitted from all following specifications, which is preferable as it increases the ability to identify a statistically meaningful result, but it does change the interpretation of the results: the estimates represent the effect of HCW mandates during a well-matched season as opposed to the effects of the mandates in expectation.

Columns 3 includes county-specific linear time trends, and the magnitudes grow substantially. Column 4 includes time-varying controls, and Column 5 includes hospital-specific linear trends in place of county trends. The event studies presented in Figure 5 illuminate the difference between specifications with and without trends. The result is substantially stronger for inpatient admissions compared to ED visits, although the patterns are similar. Diagnoses for influenza are increasing in affected hospitals relative to non-affected hospitals in years prior to mandate implementation. In the first year of implementation, however, there is a sharp decrease in the number of visits that persists in subsequent years. The inclusion of hospital-specific linear time trends and time-varying covariates appears to fully correct any pre-existing trends that were present in the model that did not include trends. Models that include trends are preferred, but in models for inpatient admissions that include

²⁹Employment in NAICS industries: 656211, 656214, 656216, 656219, 656221, 656222, 656223, 656231, 656232, 656233, 656239.

hospital-specific trends, the maximization algorithm failed to converge. For this reason, the estimates in Column 4 (with county trends instead of hospital trends) are preferred.³⁰

Also note that the comparison between Column 3 (no controls) and Column 4 (with controls) reveals only small differences in the estimates. While this is reassuring, Pei et al. (2018) note that using control variables to estimate balance regressions is a more powerful test of the identifying assumption than the more standard (and usually informal) comparison of coefficients in regressions with and without controls. I replace the outcome in Equation (5) with each of the 21 time-varying characteristics in X_{cy} (omitting the other characteristics) and find that 20 out of 21 estimates are insignificant at the 5% level, providing more support for the identifying assumption.³¹

These estimates indicate that in seasons with a well-matched vaccine, HCW vaccination mandates are associated with a statistically significant 20.1% decrease in inpatient admissions with an influenza diagnosis and a statistically insignificant 8.1% decrease in outpatient ED visits with an influenza diagnosis.³² In interpreting these results, there are at least two differences between inpatient and outpatient visits to consider. First, for inpatient visits it is possible that the individual in question was infected with influenza during their hospital stay. This means that policy-induced changes in diagnoses in an inpatient setting likely result from a combination of hospital-acquired influenza as well as influenza acquired in a non-hospital health care setting or in the community at large. For outpatient ED visits, the hospital-acquired channel is unlikely. Second, influenza may be less likely to be correctly diagnosed in an outpatient setting in which there is less time to acquire laboratory confirmation (Dugas et al., 2015). Both differences would lead to attenuated estimates for outpatient visits.

Consider next a comparison between the effects of vaccination mandates on hospital- and non hospital-acquired influenza. The final two columns of Table 7 examine the distinction between inpatient influenza diagnoses that were present on admission (POA) and those that were not (Not-POA). Diagnoses that were present on admission are more likely to represent influenza acquired outside of the hospital, whereas Not-POA diagnoses likely represent hospital-acquired infection. Influenza diagnoses are coded as present on admission for the vast majority of visits, although there is some question as to how accurate POA coding is, especially since hospitals have financial incentives to avoid coding hospital acquired conditions

³⁰The estimates are very similar using either county or hospital-specific trends. Note that there are 58 counties in California and approximately 450 hospitals (depending on the year); this means that estimating models with hospital-specific time trends requires estimating almost 400 additional parameters.

³¹p-values from this test are reported in Table B4.

³²Though the estimate for outpatient visits in Column 4 is insignificant, the estimates using slightly different specifications in Columns 3 and 5 are similar in magnitude and significant at the 10% level.

(Goldman et al., 2011).³³ The point estimate for POA influenza diagnoses is very similar to the estimate for all influenza diagnoses presented previously. Focusing on Not-POA diagnoses, however, reveals an effect twice as large in relative terms: vaccination mandates lead to an approximate 42.7% reduction in hospital-acquired influenza.

4.4.3 Additional Results & Discussion

A brief discussion of several additional results provides additional insight. Age-specific impacts of HCW mandates on both inpatient and outpatient visits are provided in Table B5. The estimates are negative for all age groups. While the differences tend not to be statistically meaningful, the magnitudes in relative terms tend to be largest for children (infants and children 1-9). This is worth noting in at least one respect: since vaccination of children should not have been affected by these policies, the implication is that much of the identified impacts operate through an externality.

Estimates for four additional outcomes that exhibit variation throughout the year are presented in Table B6, estimated using Equation (6). The first three outcomes (average length of stay, average charges, and the in-hospital death rate) are hospital-level outcomes constructed using inpatient data. The estimates indicate small and marginally significant reductions in average length of stay (1.02%) and average charges (1.96%), and no impact on in-hospital deaths during months of peak influenza activity. These outcomes are distinct from influenza-specific outcomes in that they may be viewed as measures of how well the hospital is functioning more generally. In other words, it may be that changes in these outcomes result from changes the number influenza infections, but may also reflect improvements in the health of the hospital staff, increasing the quality of their work. Column 4 presents estimates with the PI mortality rate as the outcome, and represents estimates at the county level. While it would be advantageous to examine the same outcome in Parts I and II of this paper, these estimates are underpowered to do so. None of the estimates reveal an impact of vaccination mandates on outcomes during periods in which no influenza is circulating.

Next consider an estimate of the monetary benefits of HCW vaccination. Note that the observed decreases in influenza *diagnoses* do not necessarily represent decreases in the actual number of *visits* if a substantial number of the affected visits are for other diseases that influenza infection complicated. To the extent that at least some of the observed reductions in influenza diagnoses represent a change in the number of visits, it is possible to provide a back-of-the-envelope calculation for the monetary benefit of a HCW vaccination in terms of

³³The Deficit Reduction Act of 2005 stipulated that hospitals would not receive higher payments for certain secondary conditions that were not POA. Influenza is not one of these specific conditions, but it could still be the case that hospitals develop a habit of under-reporting all hospital acquired conditions.

hospital cost savings. Focusing only on visits most likely to be avoided by the mandates, I re-estimate Equation (5) using only visits with a *primary diagnosis* for influenza (as opposed to any diagnosis) for outpatient visits, and only visits with a *primary diagnosis for influenza that was present on admission* for inpatient admissions. I consider a hypothetical statewide vaccination mandate for California (relative to no mandate) that affects all HCWs in the state. A number of assumptions are required to make this calculation, and a conservative approach is taken in applying each. I estimate that, in terms of reduced health care costs, each vaccine confers benefits of \$111.59 through reduced inpatient visits, and \$19.80 through reduced outpatient visits, for a total of \$131.40.³⁴ An estimate of the marginal social benefit would of course be larger if other outcomes such as mortality or work absences could be considered in this context as well.

Finally, it is demonstrated in the event studies that the change in influenza diagnoses coincides precisely with the timing of the mandates, suggesting that the estimates are a direct and causal result of mandate implementation. That being said, the possibility remains that health care workers simultaneously change their behavior in response to the mandates in such a way that would cloud the interpretation of the mechanism underlying the results. One possibility is that health facilities implementing vaccination mandates simultaneously take on other measures to quell the spread of disease (e.g., hand-washing or sterilization measures). This is more likely for hospital-level mandates, most of which were introduced at the height of the H1N1 pandemic. For county-level mandates, which health facilities did not directly select into, this appears to be less likely. In memos released by county public health departments that I have gathered, there is no encouragement of any behaviors other than vaccination or the requirement that workers who deny vaccination wear a surgical mask.³⁵ Table B2 demonstrates that the results are robust to the exclusion of hospitals that implemented their own mandates. Together, this suggests that a simultaneous response related to infection control on behalf of health care providers is unlikely to be a significant

³⁴This calculation is described in more detail here. First, the mean annual number of inpatient and outpatient visits for influenza in California are 4,190 and 34,221, respectively (excluding the H1N1 flu-years). The estimated coefficients using only primary diagnoses for inpatient and outpatient visits are -0.230 and -0.094, implying 964 fewer inpatient visits and 3,216 fewer outpatient visits each year. The OSHPD data include information on charges for inpatient visits, but not outpatient visits. Average charges for visits with a POA primary diagnosis for influenza are \$48,719. The Nationwide Emergency Department Sample has information on ED outpatient charges for some states, and the average charges for visits with a primary diagnosis for influenza are \$1,472. Since charges do not represent hospital costs, this figure is multiplied by the national average cost-to-charge ratio (0.507) for outpatient visits and a California-specific cost-to-charge ratio (0.288) for inpatient visits. As such, the hypothetical statewide policy achieves annual savings of \$13.52 million through reduced inpatient visits and \$2.40 million through reduced outpatient visits. To arrive at a per-vaccination figure, I divide by the number of additional vaccinations (121,154) received in a hypothetical statewide policy, discussed in Section 4.4.1.

³⁵See example memo in Figure B2.

driver of the results.

Another possibility is that when health providers are required to be vaccinated, they are more likely to encourage patient vaccination. If this is the case, an increase in non-HCW vaccination rates coinciding with the mandates would be expected. Unfortunately, this hypothesis is not testable with available data.³⁶ While these possibilities exist, the large increase in hospital worker vaccination rates in response to the mandates suggests that the intended mechanism of HCW vaccination is at work. Furthermore, even without precisely identifying all of the mechanisms at work, the reduced-form estimates of vaccination mandates on the outcomes remain valid and imply that the policies have substantial benefits.

5 Discussion and Conclusion

In this paper, I estimate the marginal social benefits of influenza vaccination for the general population (henceforth, GP) and for the population of health care workers (HCWs). Because it is not possible to use the same identification strategy and outcomes for both analyses, a direct comparison of the marginal benefits in each population requires fairly strong assumptions. I believe that such a comparison is still quite useful with this caveat in mind. To compare the benefits of vaccination in these two populations, I calculate the number of GP vaccinations or HCW vaccinations required to achieve a 1% reduction in influenza-induced mortality (for GP) or a 1% reduction in influenza-induced outpatient visits (for HCW) in a population of 100,000 individuals.³⁷ I find that 312.5 GP vaccinations are required to achieve a 1% reduction in mortality and 32.8 HCW vaccinations are required to achieve a 1% reduction in outpatient visits. Under the assumption that influenza mortality and influenza outpatient visits are proportional to each other (or to the number of influenza cases), the implication is that HCW vaccinations are nine times more effective at quelling the spread

³⁶For the period in question, data on vaccination rates (BRFSS) is not available at the county level. Furthermore, the BRFSS is not representative at the county level.

³⁷This calculation is described in more detail here. For GP, the main estimates indicate that 1,000 additional vaccinations are required to reduce annual influenza mortality by 0.249 per 100,000. To get this in relative terms, I require an estimate of annual influenza mortality, and use the nationwide 23,607 figure estimated by Thompson et al. (2010). Using the 2007 US population (to reflect the period of study in Thompson et al. (2010)), this implies 7.8 influenza deaths per 100,000 population. As such, my estimates indicate that 1,000 additional vaccinations result in a 3.2% (0.25/7.8) decrease in influenza mortality for a population of 100,000; in other terms, 312.5 vaccinations are required to achieve a 1% reduction. For HCW vaccinations, I use only outpatient visits with a primary diagnosis for influenza as this provides a lower bound (the smallest relative impacts) and I consider the same hypothetical statewide policy as in Section 4.4.3, but scaled to a population of 100,000 instead of the 2016 California population (39.3 million). For the entire state, I found that 121,154 vaccinations are required to achieve a 9.4% reduction in outpatient visits, implying that 12,888 vaccinations are required to achieve a 1% reduction. Scaling this down to a population of 100,000 implies that 32.8 HCW vaccinations are required to achieve a 1% reduction in outpatient ED visits.

of influenza in comparison to vaccinations in the general population. While this may seem large, it is not unreasonable to argue that HCWs are many times more likely to come in contact with infected individuals and much more likely to come into contact with individuals who would suffer severe consequences from being infected (e.g., a hospital visit).

For both vaccination in the general population and health care workers, I have provided policy-relevant estimates of the marginal social benefits in terms of the outcomes analyzed. For the general population, these benefits are estimated to be \$64 per vaccination in terms of reduced mortality and \$98 per vaccination in terms of work hours gained. For HCWs, the benefits are estimated to be \$131 per vaccination in terms of reduced health care costs. While the impacts of HCW vaccination on work absences and mortality could not be estimated directly in this study, one could reasonably assume that the benefits of HCW vaccination in terms of these outcomes would be at least as large as those for the GP; under this assumption, the benefit per HCW vaccination is \$293. For mortality, I find that the majority of the social benefits operate through an externality. For HCW vaccination, I do not explicitly estimate the size of the externality, but it is likely that much of the social benefit operates through the externality given that the largest relative benefits exist in a group whose vaccination status is not affected (children).

How do these benefits compare to the marginal cost of vaccination? Prosser et al. (2008) estimate that the cost of administering a vaccine (including the medicine, labor, overhead, promotion, and other expenses) ranges from \$15 in a mass vaccination setting to \$37 in a schedule doctor's office visit.³⁸ Administration costs, however, may only represent a portion of the total private costs of vaccination if there are significant non-monetary costs such as inconvenience or discomfort. Indeed, many choose not to vaccinate despite monetary costs equal to zero (influenza vaccination is covered under Medicare, and many health plans cover vaccination with zero copay). Recognizing these non-monetary costs of vaccination is critical in the development of policies that encourage influenza vaccination.

What do the estimates presented in this paper suggest for vaccination policy? The answer to this question depends on the type of policy under consideration. Let us consider two prospective vaccination policies in turn: a policy to increase vaccination in the general population and a policy to increase vaccination among health care workers.

The analysis of state-level vaccination rates is relevant to a policy that would increase vaccination among the general population by targeting those on the margin of the decision to vaccinate. Such a policy could be accomplished through a number of mechanisms: by providing monetary incentives or by reducing non-monetary costs through increasing accessibility to vaccine providers, for instance. The marginal social and externality benefits

³⁸Dollar estimates are converted to 2016\$.

estimated here suggest that vaccination policy resulting in marginal increases in the vaccination rate above current levels is beneficial so long as the marginal cost curve does not increase steeply at the current level of vaccination. While a steep increase in the cost curve is conceivable at some level of vaccination, as some individuals are opposed to vaccination on religious grounds or concerns over vaccine safety, that level is likely to be quite high as those individuals represent only a small portion of the total population (Kennedy et al., 2005). Furthermore, Bronchetti et al. (2015) show that a relatively small financial incentive can result in large increases in vaccine take-up, suggesting the existence of low-cost policies that would increase vaccination in at least some segments of the population.

For health care workers, the estimates presented here are large in comparison to administration costs. It is worth noting that many health care facilities employ mass vaccination campaigns that not only reduce the administrative costs of vaccination, but likely reduce any inconvenience costs through making vaccination highly accessible (Prosser et al., 2008; Nowalk et al., 2013). The estimates in this paper are derived from policies that *mandate* influenza vaccination, creating an extremely high cost for those choosing not to vaccinate. It is possible that other incentive-based programs could achieve a more efficient result if there are individuals for whom the marginal cost of vaccination is very high, yet still choose to vaccinate under a mandate given an even higher cost of choosing not to do so. In any case, the social benefits of health care worker vaccination estimated here are large enough to suggest that any policy increasing vaccination among health care workers would be cost-effective under reasonable assumptions about the costs.

In summary, I estimate that the social benefits of influenza vaccination are substantial and that much of the total benefits operate through externality effects. Determining the socially optimal level of vaccination depends critically on the marginal cost of vaccination – under reasonable assumptions about these costs, the results of this study indicate that policies increasing take-up of influenza vaccination in either the general population or in the population of health care workers are likely to be cost-effective.

References

- Adda, Jérôme**, “Economic activity and the spread of viral diseases: Evidence from high frequency data,” *The Quarterly Journal of Economics*, 2016, *131* (2), 891–941.
- Allison, Paul D and Richard P Waterman**, “Fixed-Effects Negative Binomial Regression Models,” *Sociological Methodology*, 2002, *32* (1), 247–265.
- Almond, Douglas**, “Is the 1918 Influenza pandemic over? Long-term effects of in utero Influenza exposure in the post-1940 US population,” *Journal of Political Economy*, 2006, *114* (4), 672–712.
- **and Bhashkar Mazumder**, “The 1918 influenza pandemic and subsequent health outcomes: an analysis of SIPP data,” *American Economic Review*, 2005, pp. 258–262.
- Althouse, Benjamin M, Theodore C Bergstrom, and Carl T Bergstrom**, “A public choice framework for controlling transmissible and evolving diseases,” *Proceedings of the National Academy of Sciences*, 2010, *107* (suppl 1), 1696–1701.
- Ashenfelter, Orley and Michael Greenstone**, “Using Mandated Speed Limits to Measure the Value of a Statistical Life,” *Journal of Political Economy*, 2004, *112* (S1), S226–S267.
- Barreca, Alan I and Jay P Shimshack**, “Absolute humidity, temperature, and influenza mortality: 30 years of county-level evidence from the United States,” *American journal of epidemiology*, 2012, *176*, 114–122.
- Belongia, Edward A, Burney A Kieke, James G Donahue, Robert T Greenlee, Amanda Balish, Angie Foust, Stephen Lindstrom, and David K Shay**, “Effectiveness of inactivated influenza vaccines varied substantially with antigenic match from the 2004–2005 season to the 2006–2007 season,” *Journal of Infectious Diseases*, 2009, *199* (2), 159–167.
- Boulier, Bryan L, Tejwant S Datta, and Robert S Goldfarb**, “Vaccination externalities,” *The BE Journal of Economic Analysis & Policy*, 2007, *7* (1).
- Brito, Dagobert L, Eytan Sheshinski, and Michael D Intriligator**, “Externalities and compulsory vaccinations,” *Journal of Public Economics*, 1991, *45* (1), 69–90.
- Bronchetti, Erin Todd, David B Huffman, and Ellen Magenheim**, “Attention, intentions, and follow-through in preventive health behavior: Field experimental evidence on flu vaccination,” *Journal of Economic Behavior & Organization*, 2015, *116*, 270–291.
- Brown, Ryan and Duncan Thomas**, “On the long term effects of the 1918 U.S. influenza pandemic,” *Working Paper*, 2016.
- Cameron, A Colin and Pravin K Trivedi**, “Count Panel Data,” *Oxford Handbook of Panel Data Econometrics*, 2013.

- and – , *Regression analysis of count data*, Vol. 53, Cambridge university press, 2013.
- Caplan, Arthur**, “Time to mandate influenza vaccination in health-care workers,” *The Lancet*, 2011, 378 (9788), 310–311.
- Carpenter, Christopher S and Emily C Lawler**, “Direct and Spillover Effects of Middle School Vaccination Requirements,” Technical Report, National Bureau of Economic Research 2017.
- Crépon, Bruno, Esther Duflo, Marc Gurgand, Roland Rathelot, and Philippe Zamora**, “Do Labor Market Policies have Displacement Effects? Evidence from a Clustered Randomized Experiment*.,” *Quarterly Journal of Economics*, 2013, 128 (2).
- Currie, Janet and Hannes Schwandt**, “Within-mother analysis of seasonal patterns in health at birth,” *Proceedings of the National Academy of Sciences*, 2013, 110 (30), 12265–12270.
- Demicheli, Vittorio, Tom Jefferson, Lubna A Al-Ansary, Eliana Ferroni, Alessandro Rivetti, and Carlo Di Pietrantonj**, “Vaccines for preventing influenza in healthy adults,” *The Cochrane Library*, 2014.
- Deschênes, Olivier, Michael Greenstone, and Joseph S Shapiro**, “Defensive investments and the demand for air quality: Evidence from the NOx budget program,” *American Economic Review*, 2017, 107 (10), 2958–89.
- Dugas, Andrea F, Alexandra Valsamakis, Mihir R Atreya, Komal Thind, Peter Alarcon Manchego, Annum Faisal, Charlotte A Gaydos, and Richard E Rothman**, “Clinical diagnosis of influenza in the ED,” *The American journal of emergency medicine*, 2015, 33 (6), 770–775.
- Dushoff, Jonathan, Joshua B Plotkin, Cecile Viboud, David JD Earn, and Lone Simonsen**, “Mortality due to Influenza in the United States An Annualized Regression Approach Using Multiple-Cause Mortality Data,” *American Journal of Epidemiology*, 2006, 163 (2), 181–187.
- Francis, Peter J**, “Dynamic epidemiology and the market for vaccinations,” *Journal of Public Economics*, 1997, 63 (3), 383–406.
- , “Optimal tax/subsidy combinations for the flu season,” *Journal of Economic Dynamics and Control*, 2004, 28 (10), 2037–2054.
- Geoffard, Pierre-Yves and Tomas Philipson**, “Disease eradication: private versus public vaccination,” *The American Economic Review*, 1997, 87 (1), 222–230.
- Goldman, L Elizabeth, Philip W Chu, Dennis Osmond, and Andrew Bindman**, “The accuracy of present-on-admission reporting in administrative data,” *Health services research*, 2011, 46 (6pt1), 1946–1962.

- Grohskopf, Lisa A, Sonja J Olsen, Leslie Z Sokolow, Joseph S Bresee, Nancy J Cox, Karen R Broder, Ruth A Karron, and Emmanuel B Walter**, “Prevention and control of seasonal influenza with vaccines: recommendations of the Advisory Committee on Immunization Practices (ACIP) United States, 2014–15 influenza season,” *MMWR Morb Mortal Wkly Rep*, 2014, *63* (32), 691–7.
- Guimaraes, Paulo**, “The fixed effects negative binomial model revisited,” *Economics Letters*, 2008, *99* (1), 63–66.
- Hausman, Jerry, Bronwyn H Hall, Zvi Griliches et al.**, “Econometric Models for Count Data with an Application to the Patents-R&D Relationship,” *Econometrica*, 1984, *52* (4), 909–38.
- Hooper, CR, A Breathnach, and R Iqbal**, “Is there a case for mandating influenza vaccination in healthcare workers?,” *Anaesthesia*, 2014, *69* (2), 95–100.
- Jefferson, Tom, Alessandro Rivetti, Carlo Di Pietrantonj, Vittorio Demicheli, and Eliana Ferroni**, “Vaccines for preventing influenza in healthy children,” *The Cochrane Library*, 2012.
- , **Carlo Di Pietrantonj, Lubna A Al-Ansary, Eliana Ferroni, Sarah Thorning, and Roger E Thomas**, “Vaccines for preventing influenza in the elderly,” *The Cochrane Library*, 2010.
- Karlsson, Martin, Therese Nilsson, and Stefan Pichler**, “The impact of the 1918 Spanish flu epidemic on economic performance in Sweden: An investigation into the consequences of an extraordinary mortality shock,” *Journal of health economics*, 2014, *36*, 1–19.
- Kelly, Elaine**, “The scourge of asian flu in utero exposure to pandemic influenza and the development of a cohort of british children,” *Journal of Human Resources*, 2011, *46* (4), 669–694.
- Kennedy, Allison M, Cedric J Brown, and Deborah A Gust**, “Vaccine beliefs of parents who oppose compulsory vaccination.,” *Public health reports*, 2005, *120* (3), 252.
- Lawler, Emily C**, “Effectiveness of vaccination recommendations versus mandates: Evidence from the hepatitis A vaccine,” *Journal of Health Economics*, 2017, *52*, 45–62.
- Lin, Ming-Jen and Elaine M Liu**, “Does in utero exposure to illness matter? The 1918 influenza epidemic in Taiwan as a natural experiment,” *Journal of Health Economics*, 2014, *37*, 152–163.
- Loeb, Mark, Margaret L Russell, Lorraine Moss, Kevin Fonseca, Julie Fox, David JD Earn, Fred Aoki, Gregory Horsman, Paul Van Caesele, Khami Chokani et al.**, “Effect of influenza vaccination of children on infection rates in Hutterite communities: a randomized trial,” *Journal of the American Medical Association*, 2010, *303* (10), 943–950.

- Manski, Charles F**, “Vaccination with Partial Knowledge of External Effectiveness,” *Proceedings of the National Academy of Sciences*, 2010, *107* (9), 3953–3960.
- , “Mandating vaccination with unknown indirect effects,” *Journal of Public Economic Theory*, 2017, *19* (3), 603–619.
- McLean, Huong Q, Mark G Thompson, Maria E Sundaram, Burney A Kieke, Manjusha Gaglani, Kempapura Murthy, Pedro A Piedra, Richard K Zimmerman, Mary Patricia Nowalk, Jonathan M Raviotta et al.**, “Influenza vaccine effectiveness in the United States during 2012–2013: variable protection by age and virus type,” *Journal of Infectious Diseases*, 2014, p. jiu647.
- Miguel, Edward and Michael Kremer**, “Worms: identifying impacts on education and health in the presence of treatment externalities,” *Econometrica*, 2004, *72* (1), 159–217.
- Molinari, Noelle-Angelique M, Ismael R Ortega-Sanchez, Mark L Messonnier, William W Thompson, Pascale M Wortley, Eric Weintraub, and Carolyn B Bridges**, “The annual impact of seasonal influenza in the US: measuring disease burden and costs,” *Vaccine*, 2007, *25* (27), 5086–5096.
- Murphy, Kevin M and Robert H Topel**, “The Value of Health and Longevity,” *Journal of Political Economy*, 2006, *114* (5).
- Neilson, William and Yancheng Xiao**, “Equilibrium vaccination patterns in incomplete and heterogeneous networks,” *European Economic Review*, 2018, *105*, 174–192.
- Neuzil, Kathleen Maletic, Beverly G Mellen, Peter F Wright, Edward F Mitchel Jr, and Marie R Griffin**, “The effect of influenza on hospitalizations, outpatient visits, and courses of antibiotics in children,” *New England Journal of Medicine*, 2000, *342* (4), 225–231.
- Nowalk, Mary Patricia, Chyongchiou Jeng Lin, Mahlon Raymund, Jamie Bialor, and Richard K Zimmerman**, “Impact of hospital policies on health care workers influenza vaccination rates,” *American journal of infection control*, 2013, *41* (8), 697–701.
- Oster, Emily**, “Does disease cause vaccination? Disease outbreaks and vaccination response,” *Journal of health economics*, 2018, *57*, 90–101.
- Osterholm, Michael T, Nicholas S Kelley, Alfred Sommer, and Edward A Belongia**, “Efficacy and effectiveness of influenza vaccines: a systematic review and meta-analysis,” *The Lancet infectious diseases*, 2012, *12* (1), 36–44.
- Pei, Zhuan, Jörn-Steffen Pischke, and Hannes Schwandt**, “Poorly measured confounders are more useful on the left than on the right,” *Journal of Business & Economic Statistics*, 2018, (just-accepted), 1–34.
- Pichler, Stefan and Nicolas R Ziebarth**, “The pros and cons of sick pay schemes: Testing for contagious presenteeism and noncontagious absenteeism behavior,” *Journal of Public Economics*, 2017, *156*, 14–33.

- Prosser, Lisa A, Megan A OBrien, Noelle-Angelique M Molinari, Katherine H Hohman, Kristin L Nichol, Mark L Messonnier, and Tracy A Lieu**, “Non-traditional settings for influenza vaccination of adults,” *Pharmacoeconomics*, 2008, *26* (2), 163–178.
- Schwandt, Hannes**, “The lasting legacy of seasonal influenza: In-utero exposure and labor market outcomes,” 2017.
- Serres, Gaston De, Danuta M Skowronski, Brian J Ward, Michael Gardam, Camille Lemieux, Annalee Yassi, David M Patrick, Mel Krajden, Mark Loeb, Peter Collignon et al.**, “Influenza vaccination of healthcare workers: critical analysis of the evidence for patient benefit underpinning policies of enforcement,” *PloS one*, 2017, *12* (1), e0163586.
- Simonsen, Lone, Robert J Taylor, Cecile Viboud, Mark A Miller, and Lisa A Jackson**, “Mortality benefits of influenza vaccination in elderly people: an ongoing controversy,” *The Lancet infectious diseases*, 2007, *7* (10), 658–666.
- Stearns, Jenna and Corey White**, “Can Paid Sick Leave Mandates Reduce Leave Taking?,” *Labour Economics*, 2018, *51*, 227–246.
- Stewart, Alexandra M**, “Mandatory vaccination of health care workers,” *New England Journal of Medicine*, 2009, *361* (21), 2015–2017.
- Stiglitz, Joseph E**, *Economics of the Public Sector.*, W.W. Norton and Company, New York, 1988.
- Stoecker, Charles, Nicholas J Sanders, and Alan Barreca**, “Success Is something to sneeze at: Influenza mortality in cities that participate in the Super Bowl,” *American Journal of Health Economics*, 2016, *2* (1), 125–143.
- Thomas, Roger E, Tom Jefferson, and Toby J Lasserson**, “Influenza vaccination for healthcare workers who care for people aged 60 or older living in long-term care institutions,” *The Cochrane Library*, 2016.
- Thompson, MG, DK Shay, H Zhou, CB Bridges, PY Cheng, E Burns, JS Bresee, NJ Cox et al.**, “Estimates of deaths associated with seasonal influenza-United States, 1976-2007.,” *Morbidity and Mortality Weekly Report*, 2010, *59* (33), 1057–1062.
- Treanor, John J, H Keipp Talbot, Suzanne E Ohmit, Laura A Coleman, Mark G Thompson, Po-Yung Cheng, Joshua G Petrie, Geraldine Lofthus, Jennifer K Meece, John V Williams et al.**, “Effectiveness of seasonal influenza vaccines in the United States during a season with circulation of all three vaccine strains,” *Clinical Infectious Diseases*, 2012, *55* (7), 951–959.
- Ward, Courtney J**, “Influenza vaccination campaigns: is an ounce of prevention worth a pound of cure?,” *American Economic Journal: Applied Economics*, 2014, *6* (1), 38–72.

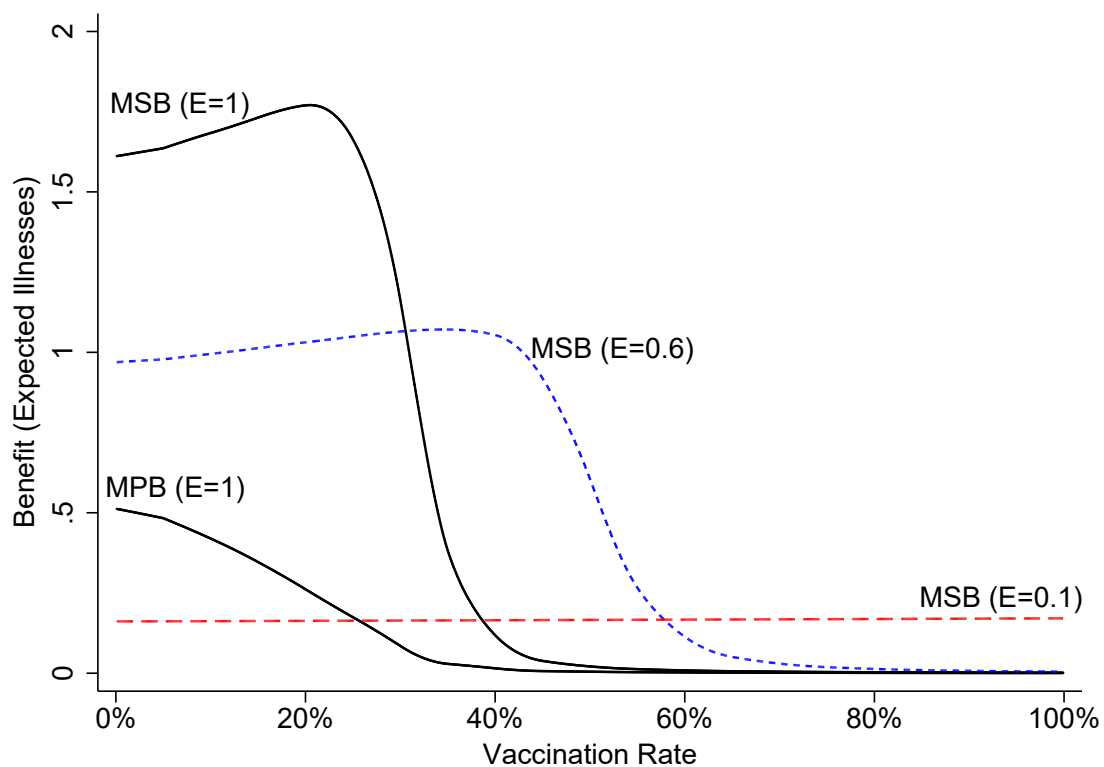
Wooldridge, Jeffrey M, “Distribution-free estimation of some nonlinear panel data models,” *Journal of Econometrics*, 1999, *90* (1), 77–97.

Zhou, Hong, William W Thompson, Cecile G Viboud, Corinne M Ringholz, Po-Yung Cheng, Claudia Steiner, Glen R Abedi, Larry J Anderson, Lynnette Brammer, and David K Shay, “Hospitalizations associated with influenza and respiratory syncytial virus in the United States, 1993–2008,” *Clinical infectious diseases*, 2012, *54* (10), 1427–1436.

Zivin, Joshua Graff and Matthew Neidell, “Environment, Health, and Human Capital,” *Journal of Economic Literature*, 2013, *51* (3), 689–730.

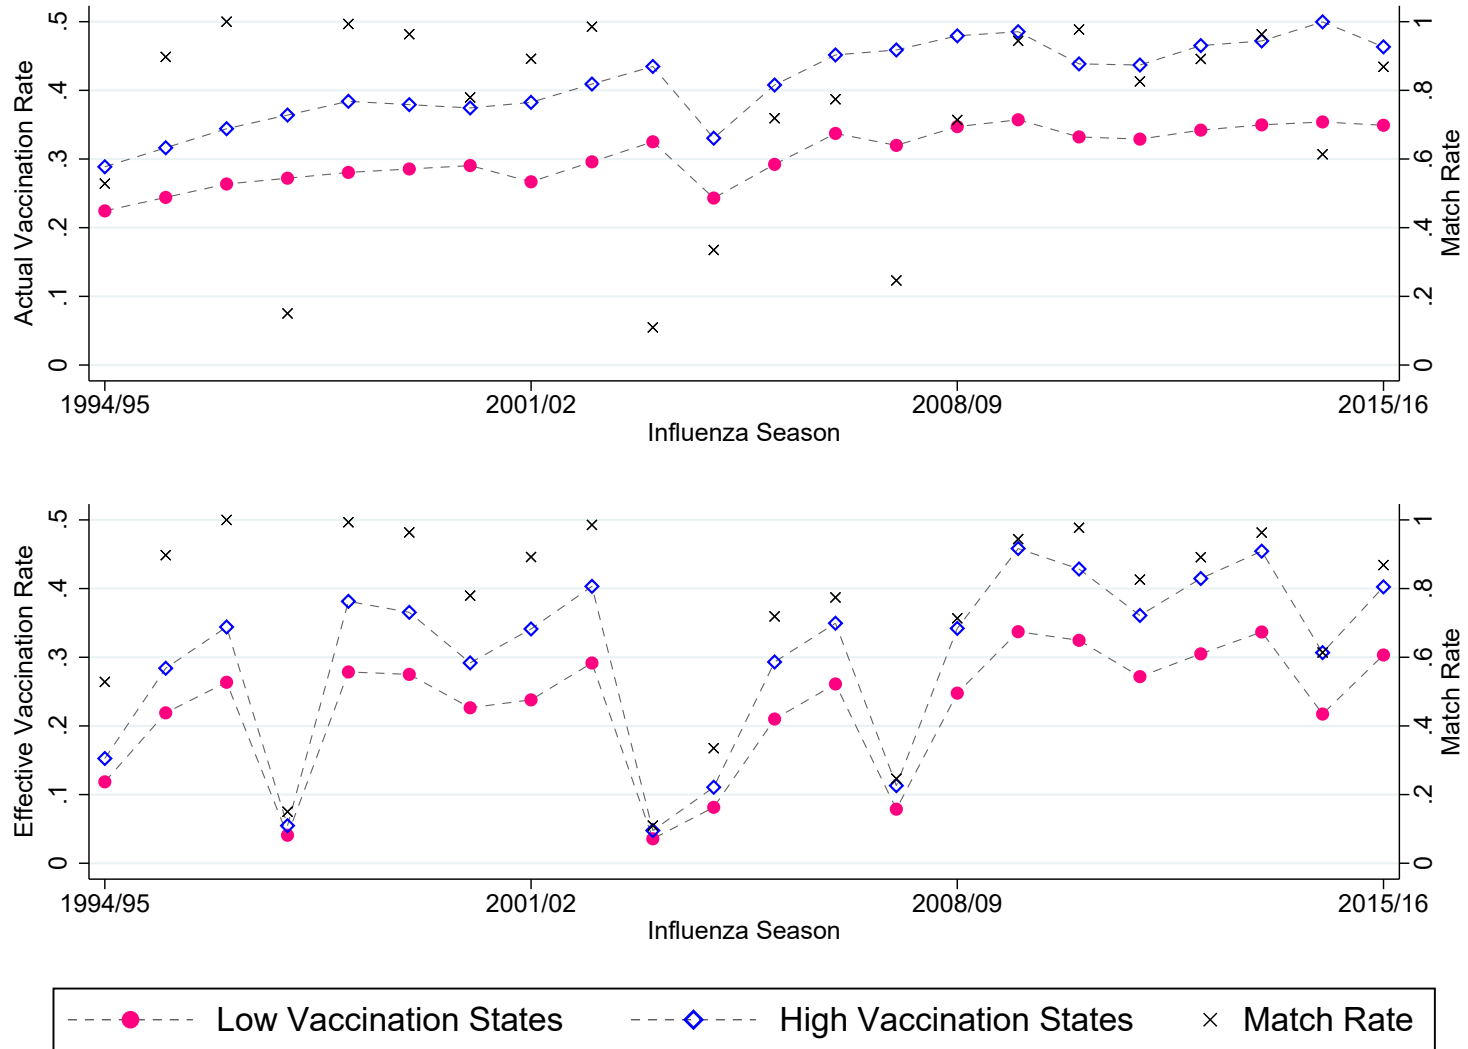
Figures & Tables

Figure 1: Marginal Benefits of Vaccination



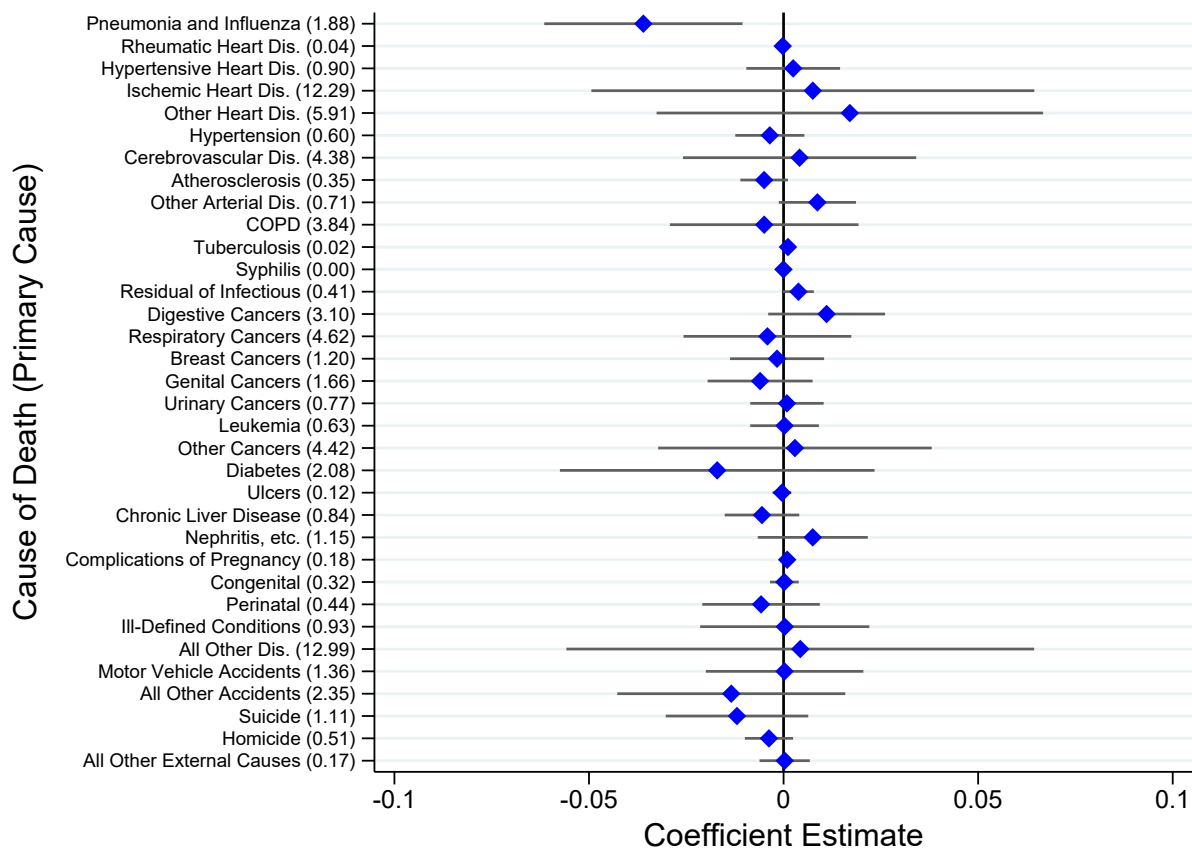
This figure presents a version of a model derived by Boulier et al. (2007) that describes the theoretical marginal social benefit (MSB) and marginal private benefit (MPB) curves for the case of influenza vaccination. These benefits are allowed to depend on vaccine efficacy. For the MSB, three levels of vaccine efficacy are presented: 100% ($E=1$), 60% ($E=0.60$), and 10% ($E=0.10$). The 10% and 60% figures represent the lowest efficacy estimates (during a low match season) and highest efficacy estimates (during a high match season) from CDC studies of vaccine efficacy for influenza seasons 2004/05-2015/16: <https://www.cdc.gov/flu/professionals/vaccination/effectiveness-studies.htm>.

Figure 2: Actual and Effective Vaccination Rates



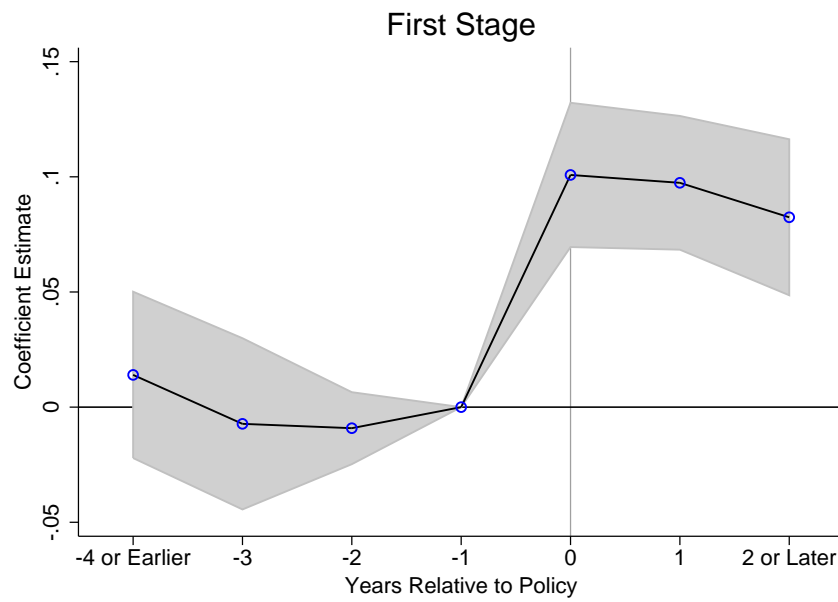
Note – There was a vaccine shortage in the 2004/05 season, accounting for the dip in vaccination rates during that season.

Figure 3: DDD Effect by Cause of Death



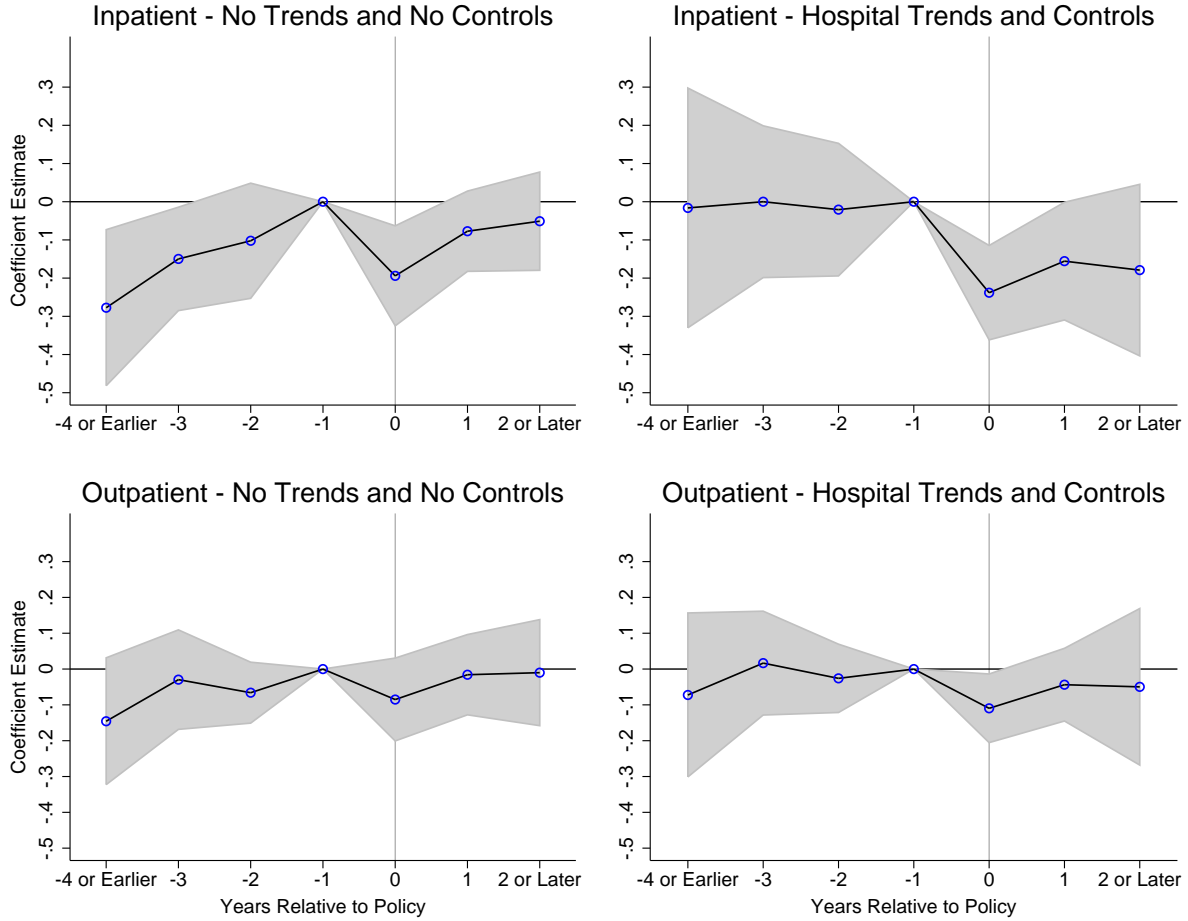
Note – Estimates represent the coefficient estimates (in levels) on the triple interaction in Equation (3) where the outcome is the mortality rate in various cause of death categories. Bars represent 95% confidence intervals. The mean monthly death rate (per 100,000 population) for each category is shown in parentheses. Cause of death is categorized by *primary* cause of death, so that each category is mutually exclusive. Cause of death categories are based on the 34-cause recode used by the NCHS for the period 1971-1998; deaths in the 1999- period were mapped from the updated 39-cause recode. Note that the mean PI death rate (1.88) is approximately one third of the death rate used in the measure for the main analysis that is based on multiple causes of death; this implies that approximately two-thirds of deaths categorized as PI in the main analysis are categorized as such based on a secondary diagnosis.

Figure 4: First Stage Event Study (HCW Vaccination Rates)



Points on this plot represent the point estimates from an event-study version of Equation (5) with HCW vaccination rates as the outcome. Shaded regions represent 95% confidence intervals. The event-study is estimated by replacing the policy indicator (Required_{hy}) with a series of variables indicating years relative to the policy: $\sum_{j=-4}^{-2} \gamma_j \text{Required}_{hyj} + \sum_{j=0}^2 \gamma_j \text{Required}_{hyj}$. The indicator representing one year prior to the policy is omitted as the reference group. “-4 or Earlier” represents four or more years prior to policy implementation; “2 or Later” represents two or more years after policy implementation. In both cases, these are aggregated because any estimates beyond this window would be identified off of a very small set of hospitals.

Figure 5: Reduced Form Event Study (Influenza Diagnoses)



Points on the plots represent the point estimates from event-study versions of Equation (5) with influenza-related inpatient admissions or outpatient ED visits as the outcome. The plot labelled “Hospital Trends” includes hospital-specific linear time trends. Shaded regions represent 95% confidence intervals. The event-study is estimated by replacing the policy indicator (Required_{hy}) with a series of variables indicating years relative to the policy: $\sum_{j=-4}^{-2} \gamma_j \text{Required}_{hyj} + \sum_{j=0}^2 \gamma_j \text{Required}_{hyj}$. The indicator representing one year prior to the policy is omitted as the reference group. “-4 or Earlier” represents four or more years prior to policy implementation; “2 or Later” represents two or more years after policy implementation. In both cases, these are aggregated because any estimates beyond this window would be identified off of a very small set of hospitals.

Table 1: Summary Statistics

National Data	Mean	SD	
Vacc. Rate	0.344	(0.063)	-
Vacc. Rate \geq 75	0.698	(0.073)	-
Vacc. Rate $<$ 75	0.315	(0.064)	-
Match Rate	0.726	(0.288)	-
Influenza Activity (1993-2015)	0.392	(0.348)	-
PI Mortality Rate	6.153	(2.037)	-
% Absent for Illness	0.025	(0.009)	-
Hours Absent for Illness	0.242	(0.117)	-
California Hospital Data	Mean	SD	# Hospitals Affected
Hospital Vacc. Rate (All Hospitals/Years)	0.801	0.142	-
Hospital Vacc. Rate (No Mandate)	0.744	(0.141)	-
Hospital Vacc. Rate (Mandate)	0.902	(0.069)	-
# Influenza Diagnoses (Inpatient)	21.81	(36.68)	-
# Influenza Diagnoses (Inpatient POA)	21.28	(35.67)	-
# Influenza Diagnoses (Inpatient Not POA)	0.51	(1.57)	-
# Influenza Diagnoses (Outpatient ED)	135.3	(170.6)	-
Average Length of Stay	5.43	(4.04)	-
Average Charges	42,294	(23,057)	-
Required 2009-10	-	-	13
Required 2010-11	-	-	18
Required 2011-12	-	-	45
Required 2012-13	-	-	116
Required 2013-14	-	-	251
Required 2014-15	-	-	335
Required 2015-16	-	-	341

Zero hospitals had vaccination mandates prior to the 2009-10 season. In the 2015-16 influenza season, the total number of hospitals in California was 450 (so that 75.8% were subject to a vaccination mandate). The total number of hospitals fluctuated between 444 (2011) and 456 (2007) over the sample period. PI refers to pneumonia/influenza; POA refers to diagnoses that are present on admission; ED refers to emergency department visits.

Table 2: Effect of the Match Rate on Vaccination Rates

	(1)	(2)
Match	-0.131 (0.295)	-3.024 (3.830)
Match \times Mean Vacc.	-	0.084 (0.112)
N	960	960

The outcome in these regressions is the vaccination rate (which varies at the state-by-flu-year level and is scaled between 0 and 100), and the regressor is the match rate (which varies only at the flu-year level). The regressions are estimated at the state-by-flu-year level. The interaction with mean vaccination rates is intended to test whether high- and low-vaccination states respond differentially to match rates. Regressions include state fixed effects and a linear time trend. Standard errors in parentheses are clustered at the state level.

Table 3: Mortality and Absences – Diff-in-Diff

Panel A: Pneumonia/Influenza (PI) Mortality Rate (per 100,000)			
	All Months	Flu Season	Non Season
Vacc \times Match	-0.043 (0.012)	-0.077 (0.023)	-0.026 (0.010)
Vacc	0.022 (0.011)	0.049 (0.022)	0.008 (0.010)
N	11,520	3,840	7,680
Panel B: Percent Absent for Illness			
	All Months	Flu Season	Non Season
Vacc \times Match	-0.00010 (0.00009)	-0.00033 (0.00014)	0.00001 (0.00009)
Vacc	0.00010 (0.00009)	0.00029 (0.00012)	0.00002 (0.00009)
N	11,520	3,840	7,680

“Flu Season” represents December-March; “Non Season” represents April-November. Standard errors in parentheses are clustered at the state level.

Table 4: Mortality and Absences – Triple Difference

Panel A: Pneumonia/Influenza (PI) Mortality Rate (per 100,000)			
	(1)	(2)	(3)
Vacc × Match × Activity	-0.087 (0.035)	-0.073 (0.031)	-0.081 (0.032)
Vacc × Match	0.010 (0.018)	-0.014 (0.013)	-0.008 (0.013)
Expected Annual Benefit (Deaths per 100,000 population)	-0.297	-0.249	-0.276
N	11,520	11,520	9,792
Panel B: Average Work Hours Absent for Illness			
	(1)	(2)	(3)
Vacc × Match × Activity	-0.00048 (0.00018)	-0.00048 (0.00016)	-0.00049 (0.00017)
Vacc × Match	0.00013 (0.00010)	0.00008 (0.00010)	0.00009 (0.00010)
Expected Annual Benefit (Hours per 100,000 population)	-4,695	-4,716	-4,874
N	11,520	11,520	9,792
Month-Year Fixed Effects	X	X	X
Weather Controls	X	X	X
State-Month Fixed Effects	-	X	X
IV	-	-	X

All regressions also include the $Vacc \times Activity$ interaction and the main effect for $Vacc$; not included are the $Match \times Activity$ interaction and the main effects for $Match$ and $Activity$ as these are absorbed by the month-year fixed effects. The “IV” specification indicates that the average vaccination rate over the three years prior is used as an instrument for the current year’s vaccination rate. Standard errors in parentheses are clustered at the state level. The “Expected Annual Benefit” for mortality is equal to $\sum_m \hat{\phi}_1 \times \overline{Match} \times Activity_m$, where $\hat{\phi}_1$ is the coefficient on the triple interaction. This measures the expected annual reduction in mortality that would be expected to result from a one percentage point increase in the vaccination rate for a population of 100,000 (i.e., 1,000 additional vaccinations). For illness absences, the “Expected Annual Benefit” is equal to $\sum_m \hat{\phi}_1 \times \overline{Match} \times Activity_m \times 17 \times (30.5/7) \times (126/323) \times 100,000$. 17 is average number of hours lost per absence. $(30.5/7)$ represents the number of weeks per month, since the coefficient measures the change in mean weekly hours lost, and $(126/323)$ represents the ratio of full time workers to the population in the U.S. Finally, the calculation for the “Expected Annual Benefit” for absences requires multiplying by 100,000 since the outcome is a proportion rather than a rate per 100,000 population. This measures the expected annual reduction in hours lost among full time workers that would be expected to result from a one percentage point increase in the vaccination rate for a population of 100,000 individuals.

Table 5: Mortality by Age & Decomposing Externality

Panel A: All-Age Vaccination Rates & Age-Specific Mortality					
	Under 1	Age 1-9	Age 10-64	Age 65-74	Age 75+
D-D-D Effect	-0.001 (0.001)	-0.001 (0.001)	-0.003 (0.005)	0.004 (0.006)	-0.073 (0.028)
N	11,520	11,520	11,520	11,520	11,520
Panel B: Age-Specific Vaccination Rates & Age-Specific Mortality					
	Under 1	Age 1-9	Age 10-64	Age 65-74	Age 75+
D-D-D Effect (75+)	-	-	-	-	0.001 (0.019)
D-D-D Effect (<75)	-	-	-	-	-0.061 (0.031)
N					11,520

Age-specific mortality rates are calculated as the number of deaths per 100,000 all-age population (i.e., the denominator is not age-specific). As such, these estimates represent an accounting of the total mortality benefits of increased vaccination – the sum of the mutually exclusive age categories equals the total effect. I report additional estimates in Table A4 in which the denominator is age-specific. In Panel A, coefficient estimates represent estimates of the triple-interaction from Equation (3). In Panel B, coefficient estimates represent estimates of the triple-interactions from Equation (4). Standard errors in parentheses are clustered at the state level.

Table 6: Monetized Benefits

Panel A: Mortality				
	Age-Adjusted VSL	Number of Deaths (Scaled Nationally)	Monetized Value (Scaled Nationally)	Value Per Vaccination
Age 75+ EPA VSL	\$975,689	807.3	\$787,673,729	\$243.79
Age 75+ AG VSL	\$256,118	807.3	\$206,764,061	\$63.99
Panel B: Work Absences				
	Median Hourly Wage	Number of Hours (Scaled Nationally)	Monetized Value (Scaled Nationally)	Value Per Vaccination
All Full-Time Workers	\$20.80	15,238,038	\$316,951,200	\$98.10

Value of a Statistical Life (VSL) estimates are generated using the EPA's figure of \$8.8 million or the estimate from Ashenfelter and Greenstone (2004) of \$2.3 million (denoted "AG"), applied to the method of Murphy and Topel (2006) to calculate age-adjusted VSL figures for each age group. Estimates correspond to a one percentage point increase in the vaccination rate, and correspond to the specification in Column 2 of Table 4. The median hourly wage is calculated as the median weekly wage in 2016 (\$832) divided by median hours worked (40) for full-time workers; these figures are derived from the Bureau of Labor Statistics <https://www.bls.gov/news.release/wkyeng.t07.htm>.

Table 7: Effects of HCW Mandates on Influenza Diagnoses

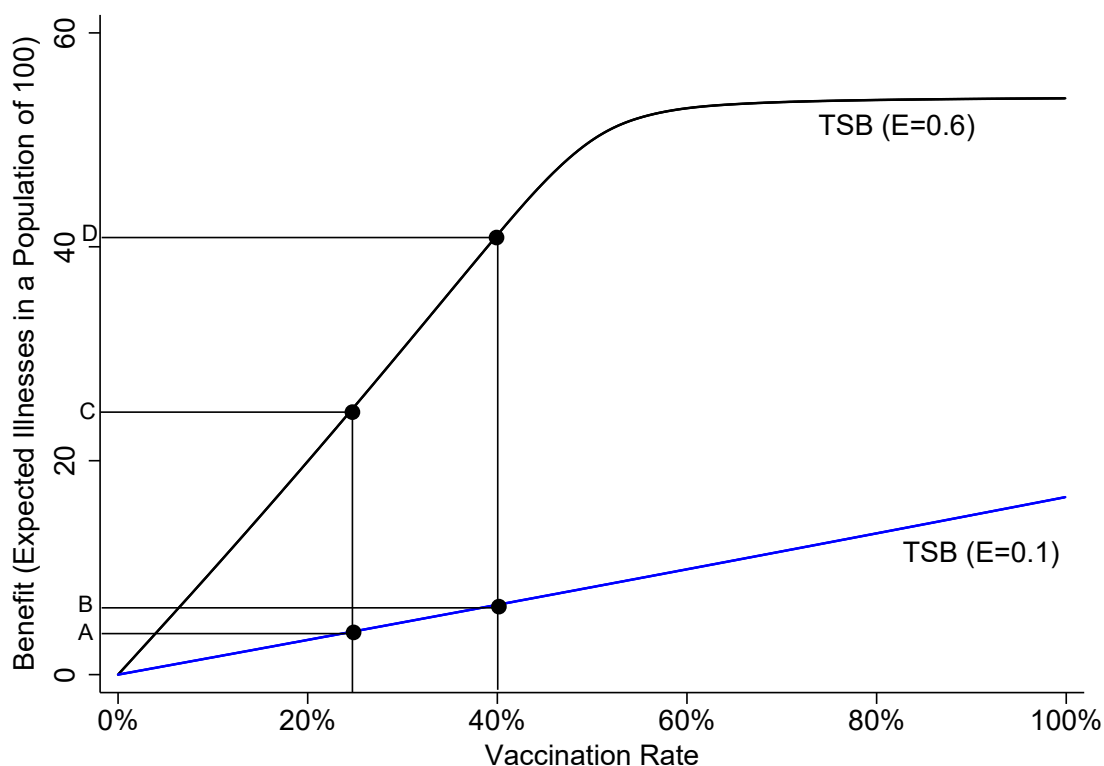
Panel A: Inpatient Admissions with Influenza Diagnosis							
	Specification Checks					Present on Admission	
	(1)	(2)	(3)	(4)	(5)	POA	Not-POA
Required	-0.0708 (0.0389)	-0.1000 (0.0459)	-0.214 (0.0687)	-0.201 (0.0775)	-0.216 (0.0783)	-0.196 (0.0757)	-0.427 (0.116)
N	3,609	3,208	3,208	3,208	3,208	3,208	3,208
Converged	Yes	Yes	Yes	Yes	No	Yes	Yes
Panel B: Outpatient ED Visits with Influenza Diagnosis							
	Specification Checks						
	(1)	(2)	(3)	(4)	(5)		
Required	-0.0168 (0.0470)	-0.0243 (0.0458)	-0.0944 (0.0544)	-0.0813 (0.0605)	-0.105 (0.0538)	-	-
N	2,700	2,400	2,400	2,400	2,400		
Converged	Yes	Yes	Yes	Yes	Yes		
Exclude 2014-15	-	X	X	X	X	X	X
County Linear Trends	-	-	X	X	-	X	X
County-Level Covariates	-	-	-	X	X	X	X
Hospital Linear Trends	-	-	-	-	X	-	-

Reported coefficient estimates are derived from negative binomial regression models, and as such the estimates can be approximately interpreted as percent changes. Regressions are estimated at the hospital-by-year level. The smaller number of observations for outpatient ED visits is due to the smaller number of emergency departments relative to inpatient hospitals. “Converged” indicates whether the maximization algorithm converged; non-convergent specifications are those that require the estimation of many parameters (i.e., hospital-specific trends) or have relatively few observations. Standard errors in parentheses are clustered at the county level.

Online Appendix

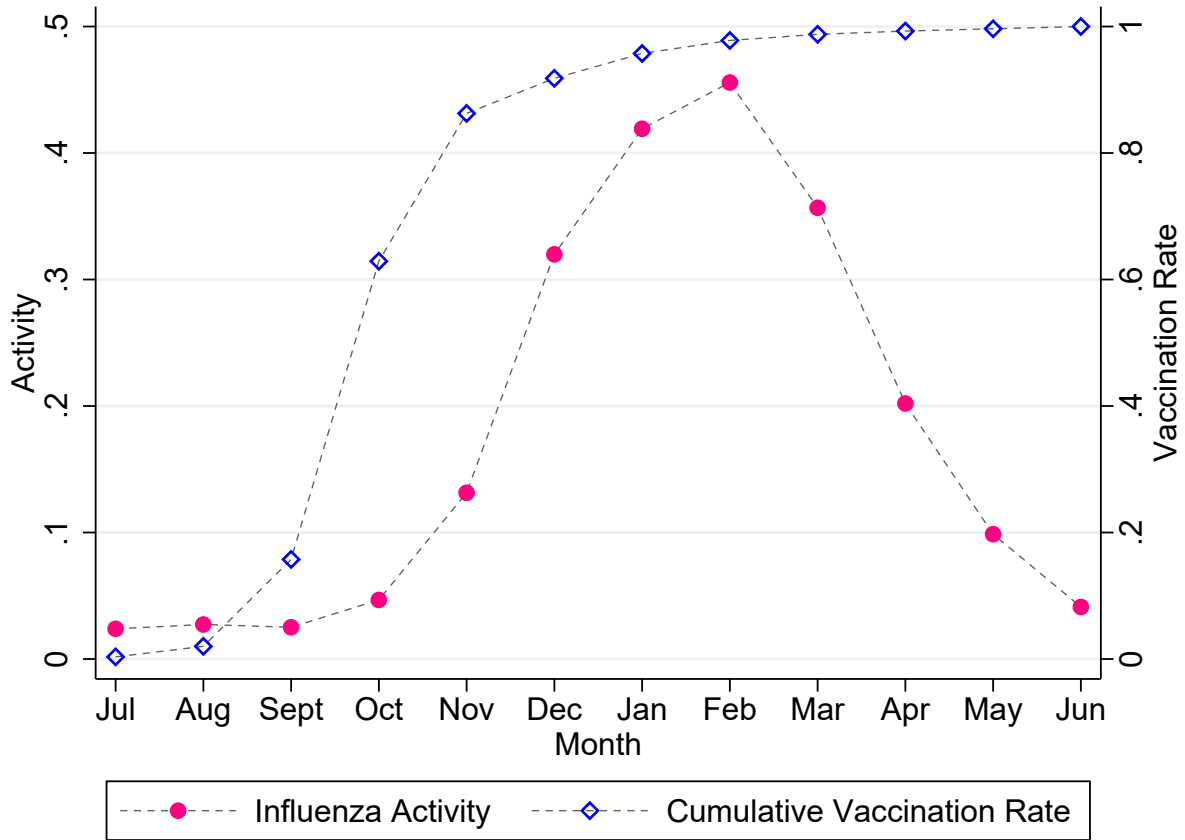
Appendix A: Part I Robustness and Additional Analysis

Figure A1: Identifying Vaccine Benefits Using Variation in Efficacy



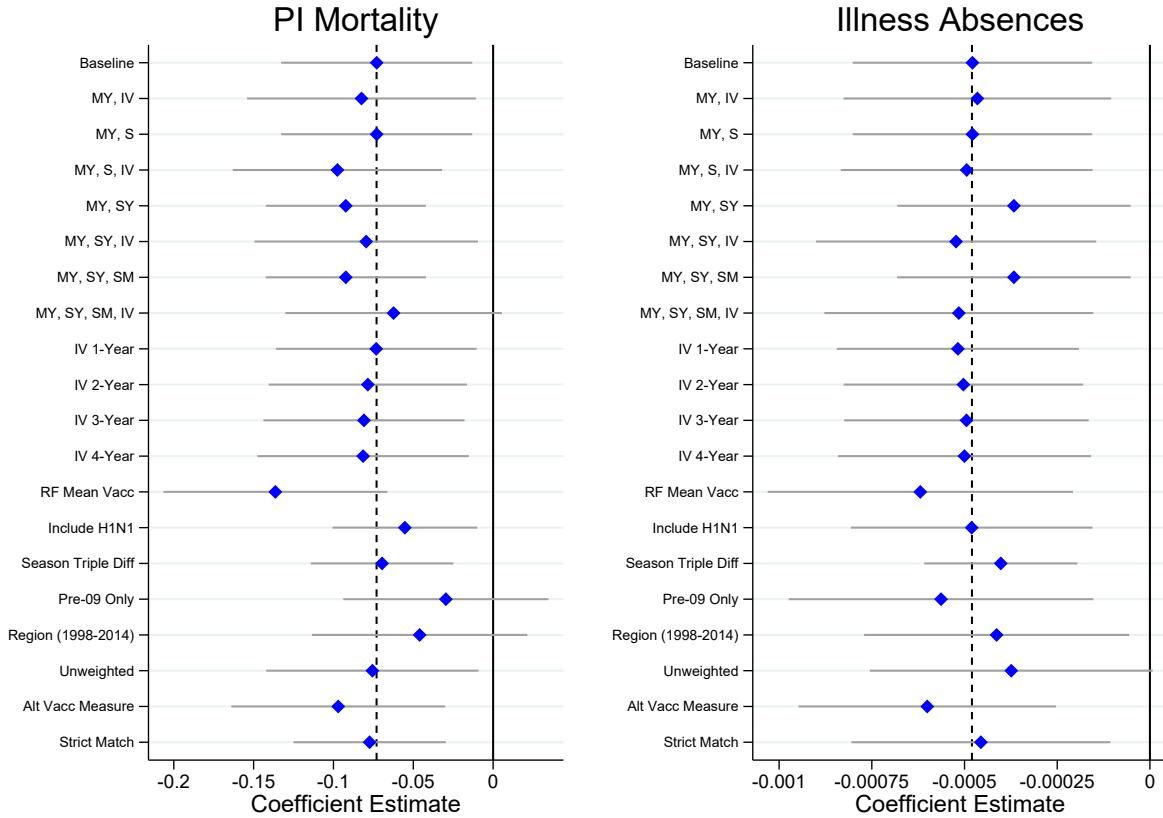
This figure presents a version of a model derived by Boulier et al. (2007) that describes the theoretical marginal social benefit (MSB) and marginal private benefit (MPB) curves for the case of influenza vaccination. Presented here are the total social benefits (i.e., the integral of the MSB presented in Figure 1). These benefits are allowed to depend on vaccine efficacy, and two levels of vaccine efficacy are presented: 60% ($E=0.60$), and 10% ($E=0.10$). The 10% and 60% figures represent the lowest efficacy estimates (during a low match season) and highest efficacy estimates (during a high match season) from CDC studies of vaccine efficacy for influenza seasons 2004/05-2015/16: <https://www.cdc.gov/flu/professionals/vaccination/effectiveness-studies.htm>. Suppose the goal is to measure the benefit of increasing vaccination rates from 25% to 40% during high efficacy ($E=60\%$) seasons: $\Delta \text{TSB} = \text{TSB}(V=40\%, E=60\%) - \text{TSB}(V=25\%, E=60\%) = D - C$. When would variation in the vaccine match rate identify this effect? If efficacy is equal to 60% during high match seasons and zero during low match seasons, then the differential effect of a high match season between low and high vaccination states identifies the same effect as increasing vaccination: $\Delta \text{TSB} = (\text{TSB}(V=40\%, E=60\%) - \text{TSB}(V=40\%, E=0\%)) - (\text{TSB}(V=25\%, E=60\%) - \text{TSB}(V=25\%, E=0\%)) = D - C$. If efficacy is not zero during low match seasons (assume $E=10\%$), then the differential effect of a high match season between low and high vaccination states identifies the following: $\Delta \text{TSB} = (\text{TSB}(V=40\%, E=60\%) - \text{TSB}(V=40\%, E=10\%)) - (\text{TSB}(V=25\%, E=60\%) - \text{TSB}(V=25\%, E=10\%)) = (D - C) + (A - B)$. So long as TSB is monotonically increasing in vaccination, the quantity $(A - B)$ is negative and the difference-in-differences identifies a lower bound on the benefit of increased vaccination.

Figure A2: Vaccination and Influenza Timing



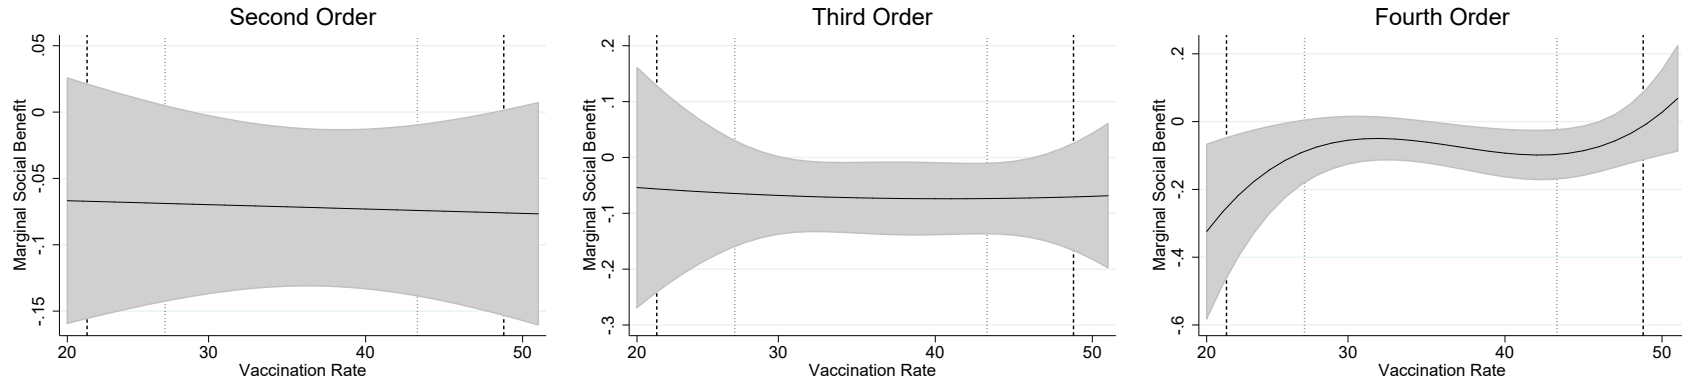
Note – This plot displays average monthly influenza activity and the average cumulative vaccination rate across years. Data on the timing of vaccination is available beginning in 2007. The year of the H1N1 influenza pandemic (2009) was excluded from the averages represented in this figure as it was a highly abnormal year in terms of the timing of both influenza activity and vaccination.

Figure A3: Robustness Analysis



Note – The “Baseline” estimate (at the top) uses the specification from Column 2 of Table 4. All other estimates are variations on the baseline model. Estimate 2-8 vary the set of controls (fixed effects) and use of the instrumental variables strategy (as in Table A1). IV denotes use of the instrumental variable, MY denotes month-year fixed effects, S denotes state fixed effects, SY denotes state-year fixed effects, and SM denotes state-month fixed effects. Estimates 9-13 vary the definition of the instrument in the instrumental variables regressions (as in Table A2): the first defines the IV as the vaccination rate in the one year prior, the second defines the IV as the average vaccination rate over the two years prior, and so on. “RF Mean Vacc” represents reduced-form estimates using a fixed vaccination rate for each state (the mean across all years); because it is a reduced-form estimate rather than an IV, the magnitude is not directly comparable. Estimates 14-20 represent a variety of other checks (as in Table A6). “Unweighted” indicates no weights are used in the regression. “Alt Vacc Measure” uses only data on vaccination that is unambiguous as to which season it applies (93/94, 95/96, 97/98, and 99/00 are omitted). “Include H1N1” includes the 08/09 and 09/10 influenza seasons. “Season Triple Diff” uses a flu-season indicator (Dec.-Mar.) in place of average monthly influenza activity. “Pre-08” uses data only from prior to the 2008/09 season. “Region (98/99-14/15)” defines the match rate at the regional level for the years in which it is possible. “Strict Match” uses a definition of the match rate that utilizes only exactly matched strains.

Figure A4: Nonlinearities



Note – These plots test for nonlinearities in the effects of state-level vaccination rates by allowing for higher order polynomials in the triple interaction specification described in Equation (3). The plots represent the marginal effects of vaccination at various vaccination rates; note that larger negative numbers imply larger social benefits. Dashed lines represent the 1st and 99th percentiles in the distribution of vaccination rates, and dotted lines represent the 10th and 90th percentiles. Shaded regions represent 95% confidence intervals.

Table A1: Mortality and Absences – Specification Checks

Panel A: PI Mortality							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
D-D-D Effect	-0.08241 (0.03655)	-0.07291 (0.03054)	-0.09753 (0.03347)	-0.09227 (0.02559)	-0.07947 (0.03570)	-0.09227 (0.02559)	-0.06234 (0.03461)
N	9,792	11,520	9,792	11,520	9,792	11,520	9,792
Panel B: Hours Absent for Illness							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
D-D-D Effect	-0.00047 (0.00018)	-0.00048 (0.00016)	-0.00049 (0.00017)	-0.00037 (0.00016)	-0.00052 (0.00019)	-0.00037 (0.00016)	-0.00052 (0.00018)
N	9,792	11,520	9,792	11,520	9,792	11,520	9,792
Month-Year Fixed Effects	X	X	X	X	X	X	X
State Fixed Effects	-	X	X	-	-	-	-
State-Year Fixed Effects	-	-	-	X	X	X	X
State-Month Fixed Effects	-	-	-	-	-	X	X
IV (t-3)	X	-	X	-	X	-	X

Estimates are reported for a variety of specifications distinct from those presented as the main results in Table 4. Standard errors in parentheses are clustered at the state level.

Table A2: Mortality and Absences - IV Specification

Panel A: PI Mortality					
	t-1	t-1 to t-2	t-1 to t-3	t-1 to t-4	Fixed (Mean)
D-D-D Effect	-0.073 (0.032)	-0.078 (0.032)	-0.081 (0.032)	-0.081 (0.034)	-0.136 (0.036)
N	10,944	10,368	9,792	9,216	11,520
Panel B: Hours Absent for Illness					
	t-1	t-1 to t-2	t-1 to t-3	t-1 to t-4	Fixed (Mean)
D-D-D Effect	-0.00052 (0.00017)	-0.00050 (0.00016)	-0.00049 (0.00017)	-0.00050 (0.00017)	-0.00062 (0.00021)
N	10,944	10,368	9,792	9,216	11,520
IV	X	X	X	X	-
Reduced Form	-	-	-	-	X

These estimates test the sensitivity of the IV estimates to the definition of the instrument for vaccination rates. Columns 1-4 use the vaccination rate in the prior year, the average vaccination rate over the prior two years, the average vaccination rate in the prior three years (the specification presented in Table 4), and the average vaccination rate in the prior four years. Column 5 uses a time-invariant mean vaccination rate (the average vaccination rate over the entire sample). Because the IV strategy requires instrumenting for three interactions plus main effect of vaccination, using a time-invariant instrument in the presence of state-by-month fixed effects means that one of the required instruments drops out. As such, these estimates are reported as reduced-form estimates (i.e., the effect of the mean vaccination rate on the outcome at time t), and the magnitudes are not directly comparable. Standard errors in parentheses are clustered at the state level.

Table A3: Mortality by Cause and Absences by Reason

Panel A: PI Mortality				
	PI Mortality	R&C Mortality	All Mortality	Non-R&C Mortality (Falsification)
D-D-D Effect	-0.073 (0.031)	-0.005 (0.065)	-0.046 (0.076)	-0.041 (0.041)
N	11,520	11,520	11,520	11,520
Mean Dep. Var.	6.15	35.52	72.16	36.63
Panel B: Hours Absent for Illness				
	% Absent - Illness	Hours Absent - Illness	% Absent - Other (Falsification)	Hours Absent - Other (Falsification)
D-D-D Effect	-0.00048 (0.00016)	-0.00517 (0.00262)	0.00022 (0.00068)	-0.00498 (0.01105)
N	11,520	11,520	11,520	11,520
Mean Dep. Var.	0.025	0.242	0.092	0.922

All estimates are from models that use the main triple-difference specification (described in column 2 of Table 4). For both mortality and work absences, the first column duplicates the main estimates as reference. “R&C” refers to respiratory & circulatory mortality, which is a level of aggregation higher than pneumonia and influenza (PI). Deaths with no respiratory or circulatory diagnosis are unlikely to be related to influenza infection (column 4 of Panel A). “Hours Absent” provides an alternate definition for work absence: the average number of hours absent rather than the proportion of workers absent. In columns 3-4, estimates are reported for non-illness work absences (e.g., vacation). Standard errors in parentheses are clustered at the state level.

Table A4: Mortality – Age Specific Mortality Rates

	Under 1	Age 1-9	Age 10-64	Age 65-74	Age 75+
D-D-D Effect	-0.065 (0.062)	-0.006 (0.006)	-0.004 (0.006)	0.080 (0.080)	-0.788 (0.413)
Mean Dep. Var.	1.27	0.12	1.17	13.02	71.33
N	11,520	11,520	11,520	11,520	11,520

in Table 5, age-specific mortality rates are calculated as the number of deaths per 100,000 *total* individuals in the population to facilitate straightforward accounting of the total mortality benefits of vaccination. Here, mortality rates are calculated as the number of deaths per 100,000 population in the relevant age group. These results confirm the main findings that the benefits are concentrated in the 75+ age group. Standard errors in parentheses are clustered at the state level.

Table A5: Mortality and Absences - Lagged Impacts

Panel A: PI Mortality					
	Baseline (One Month)	Two Month	Three Month	Four Month	Five Months
D-D-D Effect	-0.0718 (0.0308)	-0.0643 (0.0302)	-0.1318 (0.0330)	-0.1096 (0.0429)	-0.0882 (0.0603)
N	11,520	11,519	11,518	11,517	11,516
Panel B: Percent Absent for Illness					
	Baseline (One Month)	Two Month	Three Month	Four Month	Five Months
D-D-D Effect	-0.00048 (0.00016)	-0.00047 (0.00018)	-0.00010 (0.00025)	-0.00005 (0.00026)	-0.00004 (0.00035)
N	11,520	11,519	11,518	11,517	11,516

These estimates test whether the contemporaneous month is sufficient to capture the full extent of influenza-related mortality. The column labelled “Two Months” reports estimates that replicate the main estimates, but include a one month lag in the interactions that include influenza activity. The reported coefficients are the sum of the contemporaneous and lagged impact. The column labelled “Three Months” adds one and two month lags in the relevant interactions and the sum of all three coefficients are reported, and so on for columns four and five. Standard errors in parentheses are clustered at the state level.

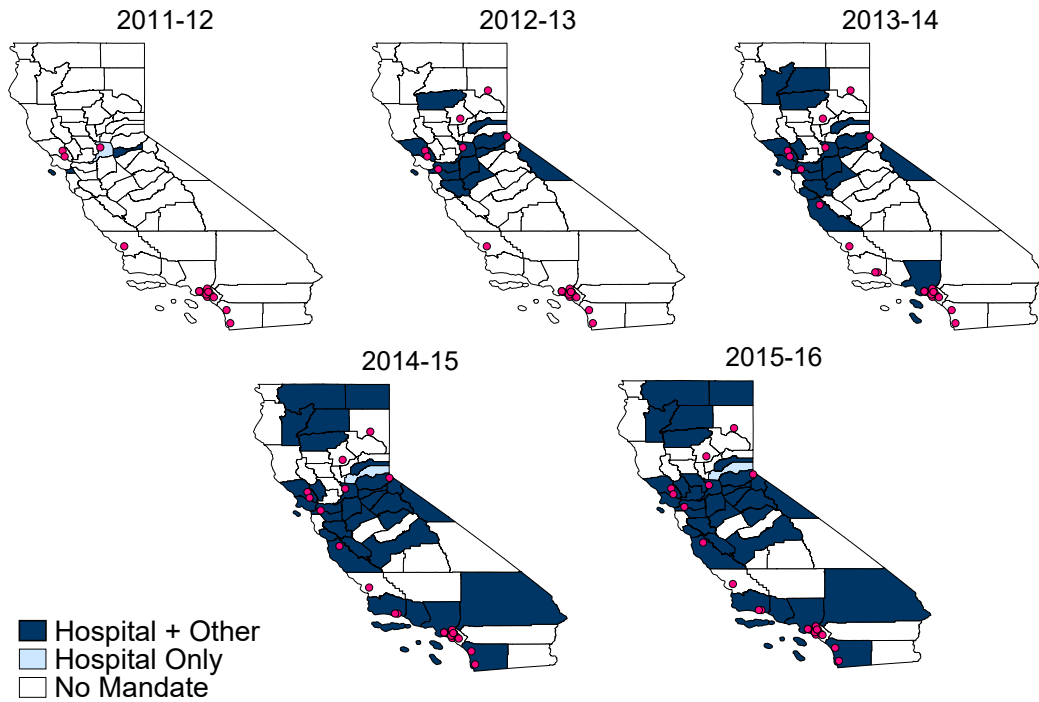
Table A6: Mortality and Absences – Robustness

Panel A: PI Mortality							
	Include H1N1	Flu Season	Pre-2009	Regional Match	Unweighted	Alt. Vacc.	Strict Match
D-D-D Effect	-0.097 (0.034)	-0.070 (0.023)	-0.030 (0.033)	-0.046 (0.034)	-0.055 (0.023)	-0.076 (0.034)	-0.077 (0.024)
Scaled Effect	-	-0.101	-	-	-	-	-
N	12,672	11,520	8,064	8,640	11,520	9,792	11,520
Panel B: Percent Absent for Illness							
	Include H1N1	Flu Season	Pre-2009	Regional Match	Unweighted	Alt. Vacc.	Strict Match
D-D-D Effect	-0.00060 (0.00018)	-0.00040 (0.00011)	-0.00056 (0.00021)	-0.00041 (0.00018)	-0.00048 (0.00017)	-0.00037 (0.00019)	-0.00046 (0.00018)
Scaled Effect	-	-0.00059	-	-	-	-	-
N	12,672	11,520	8,064	8,640	11,520	9,792	11,520

All estimates are variants on models that use the main triple-difference specification (described in column 2 of Table 4). Column 1 includes the two flu-years affected by the H1N1 pandemic. Column 2 uses an influenza season indicator in place of influenza activity in the triple difference; for this specification a “Scaled Effect” is provided for comparability with the other estimates in which the point estimate is scaled by a factor of $1/(\bar{A}_{season} - \bar{A}_{off})$, where \bar{A}_{season} and \bar{A}_{off} represent average influenza activity during influenza season and during the off-season, respectively. Column 3 represents a sample period prior to the H1N1 pandemic and two technological developments in the influenza vaccine (high-dose vaccines and the quadrivalent vaccine). Column 4 uses regional variation in the match rate (also included in this regression are the main effect for Match and the Match \times Activity interaction, which are no longer absorbed by fixed effects). Column 5 does not weight by the number of BRFSS observations. Column 6 uses an alternative calculation for vaccination rates, described in Section 3.1.2. Column 7 uses a strict definition of vaccine match, described in Section 3.1.3. Standard errors in parentheses are clustered at the state level.

Appendix B: Part II Robustness and Additional Analysis

Figure B1: California Mandates



These plots display the roll-out of influenza vaccination mandates. Circles represent policies implemented at the hospital level and shaded regions represent policies implemented at the county level. The lighter shaded regions represent county-level policies that apply only to hospitals, and the darker regions represent county-level policies that apply more broadly.

Figure B2: Example Mandate Memo – Alameda County 2013/14



Muntu Davis, MD, MPH
County Health Officer
Alameda County Public Health Department
Alameda County Health Care Services Agency
(510) 267-8000



Janet Berreman, MD, MPH
Health Officer
Health, Housing & Community Services Department
Public Health Division
(510) 981-5300



Muntu Davis, MD, MPH
County Health Officer
Alameda County Public Health Department
Alameda County Health Care Services Agency
(510) 267-8000



Janet Berreman, MD, MPH
Health Officer
Health, Housing & Community Services Department
Public Health Division
(510) 981-5300

Health Advisory
August 28, 2013

**Mandatory Influenza Vaccination or Masking of Health
Care Workers During Every Influenza Season**

Current Situation

As Health Officers for Alameda County and the City of Berkeley, we are *updating our joint Health Officer orders mandating that all licensed health care facilities in Alameda County and the City of Berkeley require their health care workers (HCWs) to receive an annual influenza vaccination or, if they decline, to wear a mask during every influenza season* (defined as November 1 to March 31 of the following year) while working in patient care areas. This order applies to all facilities regardless of documented HCW influenza vaccination rate. This updated order is effective August 28, 2013, and is ongoing and applies to each influenza season unless the order is rescinded or modified.

Background Information

Influenza infection affects 5-15% of the US population every year, leading to an estimated 3.1 million days of hospitalization and 31.4 million outpatient visits. HCWs are both at risk for influenza and can transmit the virus to their patients and coworkers. Patients in our health care facilities are particularly vulnerable to influenza. Young children, the pregnant, the elderly, and those with chronic health conditions are at greater risk for influenza-related hospitalization and death. Healthy People 2020 objectives target a 90 percent seasonal influenza vaccination rate for all health care personnel.

In your role as a health care provider and ours as Health Officers, we share common goals: reduce spread of serious diseases such as influenza, provide outstanding health care, and protect our HCWs. State law requires that general acute care hospitals and certain employers offer influenza vaccinations to employees. If employees decline vaccination, they are only required to sign a declination statement in lieu of vaccination. While compliance rates with these laws are high, actual HCW vaccination rates are not and may be below the level that will reduce the spread of infection in our health care facilities. Mandatory vaccination or masking policies have been shown to increase HCW vaccination rates to above 95%. After our first year of mandating these policies in Alameda County and Berkeley, overall HCW flu vaccination rates increased from 72% to 86% in acute care inpatient facilities.

Please distribute to all providers in your practice

Health Advisory
August 28, 2013

Our goals are to protect both patients and HCWs from influenza disease, hospitalization and death by increasing rates of influenza vaccination of HCWs and reducing HCW-to-patient transmission of influenza and vice versa.

Order:

We, as the Health Officers of Alameda County and the City of Berkeley, require that each and every licensed health care facility in Alameda County and the City of Berkeley implement a program requiring their health care workers to receive an annual influenza vaccination or, if they decline, to wear a mask during every influenza season while working in patient care areas in that health care facility.

Duration of Order

This order is ongoing and applies to each influenza season unless the order is rescinded or modified. The influenza season is defined as November 1 to March 31 of the following year. In any given year, if influenza surveillance data demonstrate an unusually late peak and continued widespread influenza activity in the spring, we may extend the period during which the masking program shall apply for that year.

Facilities Subject to the Order

This order applies to all licensed health care facilities in Alameda County and the City of Berkeley, including, but not limited to, hospitals, skilled nursing and long term care facilities, and dialysis centers. This order applies to all facilities regardless of documented HCW influenza vaccination rate.

Definition of HCWs

For the purposes of this order, "health care workers" or "HCWs" are persons, paid and unpaid, working in licensed health care settings who have direct patient contact or who work in patient care areas.

We appreciate your help and support in protecting the residents of our community. For any additional questions in Berkeley, please contact the Berkeley Public Health Division at 510-981-5300. For any additional questions elsewhere in Alameda County, please contact the Alameda County Public Health Department, Division of Communicable Disease Control and Prevention at 510-267-3230.

Please distribute to all providers in your practice

Table B1: HCW Vaccination Policy Timing

Hospital	Season	County	Season
Children’s of Orange	2009 (H1N1)	Sacramento	2011-12
Community Hospital of LB	2009 (H1N1)	San Francisco	2011-12
Hoag Hospitals	2009 (H1N1)	Alameda	2012-13
Long Beach Memorial	2009 (H1N1)	Amador	2012-13
Miller Children’s	2009 (H1N1)	Contra Costa	2012-13
Orange Coast Memorial	2009 (H1N1)	El Dorado	2012-13
Pacific Hospital of LB	2009 (H1N1)	Mono	2012-13
St. Joseph (Orange)	2009 (H1N1)	Nevada	2012-13
St. Jude (Fullerton)	2009 (H1N1)	San Joaquin	2012-13
UC Davis	2009 (H1N1)	Santa Clara	2012-13
UC Irvine	2009 (H1N1)	Stanislaus	2012-13
UC San Diego	2009 (H1N1)	Sonoma	2012-13
Saddleback Memorial	2009 (H1N1)	Tehama	2012-13
Santa Rosa Memorial	2010-11	Santa Cruz	2013-14
Sierra Vista (SLO)	2010-11	Los Angeles	2013-14
Tri-City (Oceanside)	2010-11	Marin	2013-14
Petaluma Valley Hospital	2010-11	Monterey	2013-14
Oroville Hospital	2012-13	Napa	2013-14
Banner Lassen Medical Center	2012-13	Shasta	2013-14
Barton Memorial	2012-13	Trinity	2013-14
UCSF (Children’s - Oakland)	2012-13	Alpine	2014-15
Cottage Hospitals	2013-14	Calaveras	2014-15
Salinas Valley Hospital	2013-14	Fresno	2014-15
-	-	Mariposa	2014-15
-	-	Modoc	2014-15
-	-	Placer	2014-15
-	-	San Benito	2014-15
-	-	San Bernardino	2014-15
-	-	Santa Barbara	2014-15
-	-	Siskiyou	2014-15
-	-	Tuolumne	2014-15
-	-	Ventura	2014-15
-	-	Solano	2015-16
-	-	Yolo	2015-16
-	-	Humboldt	2016-17
-	-	Mendocino	2016-17
-	-	Butte	2017-18
-	-	Merced	2017-18
-	-	San Luis Obispo	2017-18

Note: Hospitals that implemented their mandates in 2009, labelled “2009 (H1N1)”, did so in response to the H1N1 pandemic. All other mandates were implemented prior to the beginning of an influenza season.

Table B2: Effects of HCW Mandates – Robustness Checks

	Inpatient		Outpatient	
	(1)	(2)	(3)	(4)
Required	-0.199 (0.101)	-0.212 (0.0659)	-0.100 (0.0816)	-0.0833 (0.0555)
N	3,024	2,406	2,232	1,800
Converged	Yes	No	Yes	Yes
Exclude 2014-15	X	X	X	X
County Linear Trends	X	X	X	X
County-Level Covariates	X	X	X	X
Exclude Hospital Adopters	X	-	X	-
Exclude H1N1 Seasons	-	X	-	X

The outcome in first two columns is the number of inpatient influenza diagnoses and the outcome in columns 3-4 is the number of outpatient influenza diagnoses. These regressions test the sensitivity of the main result to (1) the exclusion of hospitals that adopt their own vaccination mandates (as opposed to being subject to a county-level mandate) in columns 1 and 3, and (2) the exclusion of seasons affected by the H1N1 pandemic (2008/09, 2009/10) in columns 2 and 4. Reported coefficient estimates are derived from negative binomial regression models, and as such the estimates can be approximately interpreted as percent changes. Regressions are estimated at the hospital-by-year level. Hospitals with zero age-specific influenza diagnoses in all years are automatically omitted, accounting for the difference in sample size across age groups. “Converged” indicates whether the maximization algorithm converged. Standard errors in parentheses are clustered at the county level.

Table B3: Effects of HCW Mandates on Hospital Worker Vacc. Rates

	(1)	(2)	(3)	(4)
Required	0.106 (0.0145)	0.103 (0.0169)	0.0950 (0.0113)	0.0959 (0.0113)
N	707	707	1,391	2,568
# Hospitals	101	101	202	387
90% RR All Years	X	X	-	-
90% RR All Most Years	-	-	X	-
All Hospitals	-	-	-	X
County-Level Covariates	-	X	X	X

Column 1 does not include time-varying covariates and Columns 2-4 do. Columns 2-4 represent different levels of stringency in selecting the sample for the first stage, based on data quality. Hospitals are required to report the percentage of workers vaccinated in each influenza season; poor quality data emerges when hospitals do not collect this information for every worker (the response rate for a particular hospital and influenza season may be less than 100%). Columns 1-2 reports estimates only from hospitals with a response rate (RR) of at least 90% in all years (2009/10-2015/16); Column 3 reports estimates only from hospitals with a response rate of at least 90% in all but one year (typically the first year); Column 4 reports estimates from all hospitals regardless of data quality. Standard errors in parentheses are clustered at the county level.

Table B4: HCW Balance Test

Outcome Variable	p-value ($H_0 : \pi = 0$)
Percent Uninsured	0.4643
Percent Unemployed	0.1655
Per Capita Transfers	0.9887
Per Capita Income	0.0607
Age Share 0-4	0.7256
Age Share 5-9	0.3704
Age Share 10-14	0.7074
Age Share 15-19	0.3296
Age Share 20-24	0.0100
Age Share 25-29	0.1017
Age Share 30-34	0.6061
Age Share 35-39	0.3114
Age Share 40-44	0.3473
Age Share 45-49	0.1684
Age Share 50-54	0.3203
Age Share 55-59	0.9084
Age Share 60-64	0.1641
Age Share 65-69	0.8446
Age Share 70-74	0.3577
Age Share 75+	0.6421

Note: Each row represents a county-level time-varying characteristic. Each characteristic is used as the outcome in the difference-in-differences (DD) regression model described in Equation (5), and p-values correspond to tests that the DD coefficient estimate is equal to zero. The regression models used in these tests represent the most parsimonious model, and include no time-varying covariates as explanatory variables and no county-level linear trends.

Table B5: Effects of HCW Mandates by Age

Influenza Diagnoses – Inpatient Admissions					
	Under 1	1-9	10-64	65-74	75+
Required	-0.208 (0.0866)	-0.220 (0.0979)	-0.176 (0.0912)	-0.0836 (0.124)	-0.153 (0.0692)
N	1,936	1,576	3,192	3,032	2,912
Converged	Yes	No	Yes	No	Yes
Influenza Diagnoses – Outpatient ED Visits					
	Under 1	1-9	10-64	65-74	75+
Required	-0.116 (0.0767)	-0.168 (0.0977)	-0.0767 (0.0571)	-0.112 (0.0728)	-0.0767 (0.0743)
N	2,384	2,384	2,400	2,376	2,360
Converged	Yes	Yes	Yes	Yes	No
Exclude 2014-15	X	X	X	X	X
County Linear Trends	X	X	X	X	X
County-Level Covariates	X	X	X	X	X

Reported coefficient estimates are derived from negative binomial regression models, and as such the estimates can be approximately interpreted as percent changes. Regressions are estimated at the hospital-by-year level. Hospitals with zero age-specific influenza diagnoses in all years are automatically omitted, accounting for the difference in sample size across age groups. “Converged” indicates whether the maximization algorithm converged. Standard errors in parentheses are clustered at the county level.

Table B6: Effects of HCW Mandates – Other Outcomes

	ln(Avg. Length of Stay)	ln(Avg. Charges)	In-Hospital Death Rate	PI Mortality Rate
Required×Activity	-0.0102 (0.00601)	-0.0196 (0.0118)	-0.000176 (0.000387)	-0.239 (0.283)
Required	-0.00275 (0.00337)	0.00906 (0.00940)	-0.000218 (0.000305)	0.0743 (0.0963)
Exclude 2014-15	X	X	X	X
County Linear Trends	X	X	X	X
County-Level Covariates	X	X	X	X
Hospital-Level	X	X	X	-
County-Level	-	-	-	X
<i>N</i>	33,922	30,420	33,930	5,504

The estimates presented in columns 1-3 represent average outcomes for inpatient hospital admissions at the hospital-year-month level. The estimate in column 4 uses data on mortality at the county-year-month level. The coefficient estimate for “Required×Activity” represents the impact of HCW vaccination mandates during periods with high influenza activity relative to zero activity, and the estimate for “Required” represents the impact during periods with zero influenza activity. The distributions for length of stay and charges at the micro-level (i.e., before collapsing to the hospital-year-month level) have extremely long tails. To ensure that the estimates are not driven by these outliers, I exclude micro-level observations that are above the 99th percentile of each variable’s distribution before calculating monthly averages. Additionally, charges are not reported for all inpatient visits. Some hospitals in particular consistently fail to report charges. Because these observations are unlikely to be missing randomly, I exclude hospitals that do not report charges for at least 95% of their patients over the sample period in the analysis of average charges (approximately 13% of the hospitals in the sample). This accounts for the smaller number of observations for the estimates of charges. Furthermore, length of stay is unreported for approximately 1% of admissions; averages cannot be calculated for these outcomes when all hospital-year-month outcomes are missing – this typically only occurs when there is a single observation in that cell. All models are estimated via OLS. Standard errors in parentheses are clustered at the county level.