American Economic Journal: Applied Economics 2014, 6(1): 38–72
http://dx.doi.org/10.1257/app.6.1.38

Influenza Vaccination Campaigns: Is an Ounce of Prevention Worth a Pound of Cure?[†]

By Courtney J. Ward*

This paper estimates the overall impact and externality effects of an influenza vaccination program expanding coverage outside the typical target group. Using a triple-difference design, which exploits the introduction of a broad based vaccination program in Ontario and the quality of the vaccine from year to year, I link higher vaccination to health improvements. Results indicate coverage expansion leads to large excess gains for program-regions; benefits exhibit decreasing returns corresponding to a standard model of disease dynamics; and substantial external benefits accrue to older adults. (*JEL* H75, I12, I18)

Policymakers often use evidence on the effectiveness of vaccines to set vaccine recommendations and to design health programs. However, this evidence does not generally attend to the externality effects inherent in vaccine use, which are relevant to both optimal policy response and estimation issues. To address key aspects of externalities in vaccination, this paper estimates the overall impact and externality benefits of a broad scope vaccination program expanding coverage outside the typical target group. The Ontario Universal Influenza Immunization Campaign (UIIC) was one of the first programs to both recommend and subsidize the influenza vaccine for all. Corresponding to traditional recommendations, the policy prior to 2000 included only high-risk groups, such as those over 65 and those with chronic conditions. While recommending vaccination outside traditional groups continues to be controversial, current evidence on vaccine effectiveness tends to ignore relevant features, such as the magnitude of externalities, whether benefits depend on baseline immunization levels, and how heterogeneity in efficacy can further amplify (or temper) externality effects. Furthermore, perception of these externalities may account for the broad array of

* Department of Economics, Dalhousie University, 6214 University Avenue, PO Box 15000, Halifax, NS Canada B3H 4R2 (e-mail: courtney.ward@dal.ca). I owe thanks to CEPA seminar participants at the University of Toronto; participants from numerous conference and seminar presentations; and Dwayne Benjamin, Mark Stabile, Gustavo Bobonis, Michael Baker, Robert McMillan, Philip Oreopoulos, Olivier Deschênes, and three anonymous referees. While this research draws on confidential data from Statistics Canada and the Canadian Institute for Health Information, the views expressed herein represent only those of the author. The author acknowledges financial support from the Social Sciences and Humanities Research Council of Canada.

[†] Go to http://dx.doi.org/10.1257/app.6.1.38 to visit the article page for additional materials and author disclosure statement(s).

recent policy responses in the face of looming disease epidemics, which often do not accord with existing recommendations.[1]

Such concerns over impending seasonal epidemics are justified. Influenza is the sixth leading cause of death among adults, killing as many as breast cancer and three times as many as HIV/AIDS in the United States annually (Thompson et al. 2003, 2004; Poland, Tosh, and Jacobson 2005; World Health Organization (WHO) 2006). Given medical care costs alone, it is not surprising that public health officials emphasize vaccination as a way to mitigate these costs and, hence, focus these efforts on the elderly, where mortality risk is highest. At the same time, children are the highest risk for infection (Fiore et al. 2010), and while influenza is linked to mortality more directly, early infection may lead to longer term disability and worse educational outcomes later in life (Almond and Mazumder 2005; Almond 2006).

By focusing on a broad-based program expanding coverage to children and younger adults, this study addresses the limitation of previous evidence by assessing excess gains from vaccination, showing that higher efficacy rates in the young can lead to externality benefits for the old, and showing that overall gains decline to exhaustion depending on baseline immunization levels. Since a simple before and after comparison for UIIC regions may incorrectly attribute all changes in health to the program and even conventional difference-in-differences comparisons with non-UIIC regions may be confounded with differential trends, the analysis, instead, relies on variation in the quality of the influenza vaccine. Because this variation is based on the degree to which the vaccine relates to the influenza virus each year, it can be considered as exogenous to the decision to vaccinate. For instance, while the genetic composition of influenza is constantly changing, the composition of the vaccine is predetermined and fixed for each flu season and across North America. This implies that if there are benefits from vaccination, UIIC regions will have much to gain when the vaccine is a good match, whereas they will have little to gain when it is not.

The analysis proceeds by combining survey and administrative data to construct weekly measures of work absences, hospital admissions, and influenza surveillance for sub-provincial regions over an 11-year period, which allows for flexible specifications including season, month, and region fixed effects. Using these data, the analysis informs five main sets of results. First, the overall program effect is sizable. During the epidemic period and when the vaccine is well matched, there is a 48 percent decrease

[1] For instance, the recommendations of the US Advisory Committee on Immunization Practices (ACIP) prior to 2008 were set for the standard "high-risk" groups (Fiore et al. 2007). In 2008, the ACIP expanded recommendations to include children under 18, but did not include ages 19–64, stating that "economic analysis among adults aged <65 have reported mixed results regarding influenza vaccination" (Fiore et al. 2008). Meanwhile, the threat of influenza H1N1 in 2009 led to rapid and varying policy responses. At the local level, some states and/or employers went as far as mandating vaccination (http://www.cdc.gov/h1n1flu/vaccination/acip.htm), and at the time, the official position of the ACIP was: "that as many people as possible receive 2009 H1N1 vaccine as quickly as possible" (http://www.cdc.gov/h1n1flu/cdcresponse.htm). Following the 2009 H1N1 pandemic response, the ACIP maintains that evidence for the cost effectiveness of the seasonal influenza vaccine for adults aged <65 remains mixed, but it has adopted the policy of recommending the seasonal influenza vaccine to all groups older than six months (Fiore et al. 2010; CDC 2011, 2012). In Canada, the National Advisory Committee on Immunization (NACI) currently recommends the vaccine to typical target groups, whereas remaining groups ages 5–64 are only "encouraged" to vaccinate pending further review (NACI 2012). Meanwhile, five provinces currently operate universal coverage programs for the influenza vaccine, while the other provinces retain targeted coverage. In the United Kingdom, the National Health Service continues to limit vaccine coverage to the standard target groups (http://www.nhs.uk/Conditions/vaccinations/Pages/who-should-have-flu-vaccine.aspx).

in influenza-pneumonia admissions, a 14 percent decrease in work absenteeism, and a decrease in the rate of influenza-pneumonia deaths of 9 per 100,000. On the other hand, there is little effect for any of these measures in the *nonepidemic* period and little effect for nonrespiratory admissions in *any* period. Second, for those over 65, coverage status and, likewise, relative vaccination remained unchanged in UIIC regions post program. Even so, this group experienced substantial decreases in respiratory admissions and death. Third, estimated excess benefits for all age groups are large. A benchmark estimate of the decrease in infections under a high match is 92 percent while the size of the decrease in the unvaccinated population is only 11 percent. Fourth, corresponding to the predictions of a standard theoretical model of influenza dynamics, the results show that these gains would fall to zero at vaccination rates above the mid-30's (this result considers the case of a well matched vaccine which occurs in one out of every 2.2 seasons). Finally, the implied aggregate benefits of the program, including spillover benefits to older adults, are substantial relative to program costs. Considering only hospitalization and productivity losses, the impact of the vaccine campaign translates into best-case scenario cost savings of $241 million in a high match season or expected savings of $171 million in an average match season. Meanwhile program costs are about $33 million annually.

The rest of the paper proceeds as follows. Section I provides background information on the influenza virus and influenza vaccines. Section II outlines the identification strategy. Section III describes the data and presents descriptive statistics. Section IV presents results and interpretation and Section V concludes.

I. Background

A. *Influenza*, *Vaccine Efficacy*, *and Vaccination Programs*

Each year almost 1 out of every 10,000 Americans die from influenza and its complications. For those over 65, this number is 1 in 20. For those aged 18 to 64, influenza is responsible for an average 75 million days of work absenteeism, and 22 million health care provider visits annually (Benson and Marano 1998; Thompson, et al. 2003, 2004; Poland, Tosh, and Jacobson 2005; Fiore et al. 2010). Research has also linked early influenza infection to added long-term effects. For instance, children in utero during the 1918 H1N1 influenza pandemic displayed increased rates of physical disability, and decreases in income and educational attainment later in life (Almond and Mazumder 2005; Almond 2006).

The disease itself, is contagious through droplet spread, with contagion risk beginning prior to the onset of symptoms and continuing for a number of days after recovery.[2] The leading defense against influenza is vaccination, which provides

[2] The acute effects of influenza can last from 3 to 14 days (Fiore et al. 2010; PHAC 2006). In addition, the virus can stay virulent on surfaces for a varying length of time. At body temperature, the virus is usually inactivated in less than a week, whereas in cool dry temperatures the virus can last considerably longer (Zhang et al. 2006). This is, in part, the reason why seasonal epidemics appear during winter months (WHO 2003). Further information on influenza and vaccination is available from the US Centers for Disease Control and Prevention (http://www.cdc.gov/flu/about/disease/index.htm) or the Canadian Immunization Guide published by the Public Health Agency of Canada (http://www.phac-aspc.gc.ca/publicat/cig-gci/index-eng.php).

protection from infection by triggering an antibody producing immune response to targeted influenza strains. This means that vaccine efficacy depends on the immune response of the recipient, which is compromised if antibody levels are too low (Fiore et al. 2010). For instance, systematic review indicates efficacy rates in healthy adults of 80 percent, while the results for children and older ages are lower: 62 percent for children, and 58 percent for adults over 65 (Govaert et al. 1994; Jefferson et al. 2007; Manzoli et al. 2007). Furthermore, due to antigenic drift in the influenza virus, seasonal protection also depends on the match of vaccine strains to circulating strains of influenza.[3] For instance, in a systematic review, Jefferson et al. (2007) note lower effectiveness in studies taking place in seasons where the vaccine is not well matched to circulating strains.

Due to constant genetic change in the influenza virus and the relevance of vaccine match to efficacy, the vaccine is reformulated each year to account for changes in the antigenic composition of influenza strains. WHO monitors circulating influenza viruses and in early spring writes the annual vaccine recipe. The vaccine includes two subtypes of influenza A (H3N2 and H1N1) and one influenza B virus. Usually, one or two of the three virus strains in the vaccine are changed each year and the prescription is identical across the North American continent (WHO 2003). In Canada, Health Canada licenses the newly formulated inactivated vaccine each year, and once licensed Public Works and Government Services Canada purchases influenza vaccines on behalf of the provinces for distribution in late October early November.[4] The turn around period from the yearly WHO recommendation to availability is around six to eight months depending on manufacturing conditions (Health Canada 2007).

Beginning in the early 1990s, all Canadian provinces developed programs covering the cost of vaccination for specific "high-risk" groups. The standard covered group included recipients less than 24 months, older than 65 years, health care support staff, residents in care homes, and those with specific chronic conditions (Johansen et al. 2004).[5] Aside from minor changes to provincial programs, and with the exception of Ontario, vaccine targeting remained the same across all provinces for the next two decades. In July of 2000, the Government of Ontario announced the extension of vaccination coverage to all residents of the province through the Universal Influenza Immunization Campaign (MOHLC 2000). The campaign was initiated as part of a larger ten point plan to address emergency room overcrowding, which highlighted general resource constraints rather than, specifically, a high influenza burden (for instance, influenza admission rates and surveillance rates were marginally lower in Ontario relative to other provinces) (Kurji 2004). In its first

[3] Specifically, the genetic structure of the virus is constantly changing over time through point mutations, which leads to different strains of the virus that are genetically differentiated on the basis of surface antigens. Antigenic drift in the influenza virus can mean, depending on antigenic changes, minimal cross-immunity for any new strain (PHAC 2006).

[4] Administrative data from Ontario OHIP physician billings show that the majority of yearly vaccination takes place before December (Kwong and Manuel 2007). Similarly, survey data from the US (NCHS 2008) show that 91 percent of vaccinated respondents had received the current influenza vaccine by December.

[5] Conditions include cardiac or pulmonary disease, asthma, diabetes, renal disease, liver disease, anaemia, HIV, and cancer.

year, the program cost $31 million with 2.1 million additional vaccines delivered (Kurji 2004).[6]

B. *Externalities and the Literature on Vaccine Effectiveness*

Policies for annual universal vaccination, such as UIIC, go against traditional vaccine recommendations. For instance, based on meta-analysis of several randomized evaluations, Demicheli (2001) claims that benefits from vaccinating healthy adults are small and "at odds with the conclusions reported in previous meta-analysis of evidence for the effect of immunization on elderly people, which showed greater clinical effectiveness, thus supporting the present worldwide policy of vaccinating only elderly people and other high-risk groups." Further to this, more recent surveys indicate that evidence on the effectiveness of vaccinating healthy adults continues to be mixed with no uniform prescription on the validity of recommending vaccination for these groups (Nichol 2001, 2003, 2008; Fiore et al. 2007, 2008, 2009, 2010; CDC 2011, 2012).

However, externalities in vaccination can complicate this evidence in at least two important ways: first in estimation and second in socially optimal policy setting. For instance, vaccine recommendations are largely based on evidence from randomized control trials (see Fiore et al. 2010 for a complete review), which may yield undervalued results depending on the degree of external spillovers in each context. Specifically, studies often randomize vaccination within a chosen sample, compare outcomes across treatment and control, but fail to deal with external benefits for the untreated group that accrue from vaccination of the treated (which may be large if the sample is chosen from, for instance, a specific locale or workplace).[7] Furthermore, population externalities can lead to differential returns to treatment conditional on underlying vaccination characteristics, which can yield diverse results across study context depending on the degree of local vaccination and compositional take-up.[8] For social planners, not only are estimates of vaccine effectiveness potentially undervalued due to externalities, the evidence provides little direction for the size of externalities or how they relate to issues of compositional take-up or differential returns.[9] On the other hand, information on these relationships is valuable for

[6] Overall, the total program cost is a small portion of the health care budget of more than $30 billion annually. Furthermore, program funding is independent of hospital and physician budgets, which are determined separately through funding formulas and negotiations with the Ontario Medical Association. Since coincident policies for general emergency room overcrowding likely have implications for hospital admission, the analysis on hospitalizations controls for changes in resources devoted to other aspects of the public health care system. As discussed below and formally in Section III, there is still cause for concern in attributing simple post-UIIC difference-in-differences estimates as due to the UIIC program alone.

[7] This difficulty has been shown in other contexts. See, for instance Philipson (2000a) for a general discussion and Miguel and Kremer (2004) for a discussion within the context of intestinal helminthes and deworming interventions.

[8] Externalities are not the only challenge facing evidence from randomized evaluations. Since these studies typically isolate only one flu season, they cannot adequately account for the role of the influenza vaccine match, which impacts the estimated benefits of vaccination systematically.

[9] Furthermore, since composition and level of vaccination in the underlying population varies across study context and season, this will have implications for estimated effectiveness even where within-study estimation spillovers are not an issue. For example, each evaluation typically involves a small sample, representing a specific group, within a specific context, observed over one flu season; it then remains unclear how well these results can be combined to inform a broad-based vaccination policy for a more generalized population. Take the following case

governments, firms, or health agencies wanting to design policies that exhaust external net-benefits of prevention, and as the analysis will show, externalities can have meaningful implications for both estimation and policy design.[10]

II. Identification Strategy

The purpose of the empirical work is to study the links between immunization and health by identifying the impact of vaccination as delivered through a broad scope program. To start, the following model links underlying vaccination to health (broadly defined) with linear and homogenous effects in vaccination:[11]

$$h_{jt} = \alpha + \beta m_{jt} v_{jt} + \varepsilon_{jt}. \quad (1)$$

Here, h_{jt} is a health measure for region j at time t; m_{jt} reflects vaccine quality and is measured as the match rate between the vaccine and circulating influenza strains; v_{jt} is the vaccination rate; and ε_{jt} is the error term. Both own and external effects of vaccination are subsumed in β, which represents the total effect of changes in average vaccination, scaled by vaccine quality, in region j at time t.[12]

Because β captures both the individual and external effects of vaccination, its magnitude is of interest to policymakers hoping to recoup societal benefits through increased vaccination rates. However, estimation of β using variation in vaccination rates presents a number of potential challenges. First, while higher vaccination rates may lead to better health, it may also be the case that poor health leads to higher vaccination rates. This association between poor health and take-up of vaccination may mitigate the estimated relationship described by β

abstracting from all but the issue of compositional take-up, and consider a randomized evaluation for two groups: those under 65 and those over 65. First, on a theoretical basis, vaccination for those under 65 versus those over 65 could lead to knock on compositional effects if different efficacy rates among the young impact externality benefits for the old (e.g., policies targeting the younger group may lead to higher gains for *both* groups versus the alternative scenario). However, where evaluations are context specific, considered in isolation, and do not reflect on these spillovers, they will remain silent on this issue. Second, within age group, randomized versus optional take-up may lead to different effects if factors (such as health) are related to decisions to vaccinate and to costs of infection (e.g., under optional take-up, those with higher expected health gains may be the first to vaccinate, and while selection of this sort would be desirable to the policymaker, each evaluation is unable to reflect on this effect). In general, assessments of evidence from randomized evaluation for specific groups will not be able to address these larger issues of compositional take-up, knock on effects, or the related issue of differential returns to vaccination.

[10] Policies designed to recoup social benefits could be warranted if private action fails to exhaust all external net benefits of prevention. Furthermore, for the case of vaccination, this gap may be even larger if, for instance, health coverage is offered for treatment of disease versus disease prevention, thereby distorting the relative price for preventative care and lessening the incentive to prevent illness. However, even where health authorities consider only private benefits in setting vaccine recommendations, the estimation issues highlighted above may still confound these objectives.

[11] While the assumptions of linearity and homogeneity are likely too restrictive in the current context, the following is used as a baseline model to describe core estimation issues before exploring the implications of nonlinearity and heterogeneity.

[12] Here, β is a function of the individual and external effects of vaccination obtained from an individual model:

$$h_{ijt} = \alpha + \delta m_{jt} v_{ijt} + \gamma m_{jt} v_{(i)jt} \varepsilon_{jt}.$$

In the individual model, the coefficient δ captures the individual effect of vaccination (denoted v_{ijt}) from person i's vaccination decision, and the coefficient γ captures the effect of average vaccination excluding person i (denoted $v_{(i)jt}$). In other words, it is the external effect of vaccination on the illness of person i arising from the vaccination behavior of others. The *total effect* of vaccination is recoverable when averaging the individual model over region and time.

and would bias estimates of the effect of vaccination downward.[13] Additionally, there may be other factors that contribute to region-by-year variation in vaccination rates. For example, variation in incidence of other infectious diseases (colds, other respiratory viruses, et cetera) may be correlated with selection into vaccination and also associated with health outcomes. In this case, the positive association between vaccination and health would be mitigated by this correlated effect. Lastly, vaccination, itself, may be associated with other behaviors that affect health. For instance, hand washing or other prevention methods may increase after receiving the vaccine. If, for instance, a higher vaccination rate is associated with increased exposure to prevention information and increased prevention behaviors, then rates of illness may be lower regardless of receipt of the shot. If such an effect is not delivered through vaccination itself, but instead through an advertising effect regarding the upcoming flu season, this has separate implications for policy, especially considering the inoculation costs of vaccination, in addition to the costs of time loss and personal discomfort.

In the present study, there are several factors that contribute to identification of the effect of a broad scope immunization program. The first factor is time-region variation in a program that provides vaccines free of charge to all age groups. The second factor is that in different years and across provinces, there are different degrees of match between the vaccine and influenza. Furthermore, these mismatches are determined by random mutations in influenza as they relate to yearly vaccine content, which is predetermined and fixed across North America and over each year. This means that, unknown to the recipient at the time of vaccination, the vaccine may offer a high degree of protection or it may offer a marginal degree of protection. Accordingly, areas with higher levels of vaccination will experience greater benefits if the vaccine is a good match and smaller benefits when it is not.

These elements combined form the basis for disentangling differential changes in health from the causal impact of the immunization program. For instance, aspects such as underlying differences in labor conditions, the supply of hospital services, and sentiment toward prevention behavior may evolve differentially over time and across region. Without considering the impact of these alternate, often unobservable, drivers of health, estimates of the value of immunization programming can be either undervalued or overvalued depending on these underlying factors.[14] Given that results are sensitive to these underlying factors, the methodology

[13] In fact, the data show that this is a probable concern. Comparing vaccinated and unvaccinated groups; the data show a counterintuitive connection between vaccination and health during the nonepidemic period where the vaccine is unlikely to causally impact health. Specifically, in the summer months where influenza is not in circulation, the rate of recent short-term illness is 16 percent higher for vaccinated versus unvaccinated individuals. This finding likely demonstrates selection into vaccination where those in poorer health or with higher infection probabilities are more likely to vaccinate, but are also more likely to experience negative health shocks.

[14] For instance, Groll and Thompson (2006) find that there is a small relative *increase* in surveillance counts of laboratory-confirmed influenza for Ontario compared to other provinces at the introduction of the universal program. This result is likely explained by the steeper upward trend in testing in Ontario relative to the other provinces. Meanwhile, Kwong (2005), using time variation in hospital counts, finds little relative impact for influenza hospitalization for Ontario post-program. Alternatively, using a similar methodology, but extending the yearly time series, Kwong et al. (2008) find relative decreases in death, hospitalization, and emergency room visits at the introduction of the campaign. These results could be explained by differential trends in health resources that widen over time. For instance, over this period, the number of *nonrespiratory* admissions declined by 1.5 percent per year in Ontario, but did not change for the other provinces.

used here explicitly controls for agnostic differences in health accruing to UIIC regions post-program and identifies the impact of vaccination by using exogenous variation in vaccine quality:

$$(2) \quad y_{jt} = \beta_1(UIIC_p \times Post_y \times m_{py}) + \beta_2(UIIC_p \times Post_y) + \beta_3(Post_y \times m_{py}) + \beta_4(UIIC_p \times m_{py}) + \mathbf{X}'_{jt}\Pi + u_{jt}$$

where y_{jt} is an illness outcome for region j at week t.[15] On the right hand side $UIIC_p$ is an indicator for regions in Ontario; $Post_y$ is an indicator for seasons after 2000; and m_{py} is the clinical match rate (ratio of vaccine-matched strains to total number of tested strains) in province p in season y.[16] The vector $\mathbf{X}_{jt}$ is a set of controls for: public health expenditures on health care (hospitals, capital, physicians and other health professionals); diagnosis specific coding changes from ICD9 to ICD10; the match rate in levels; and region, season and month fixed effects.[17] The variable u_{jt} is an error term clustered at the region level to account for both serial correlation and correlated random shocks within region and season.

Inclusion of region effects captures fixed features among economic-regions, and will account for unobservable region differences that are common across all seasons. Similarly, by controlling for season fixed effects, the model accounts for any fixed differences across seasons common to all regions. Finally, the coefficient β_2 directly captures unobservable post-period differences in health accruing to UIIC regions.

The model also includes careful accounting for the baseline effects of the vaccine match to allow for the possibility that, among strains of differing severity, the vaccine selection process gives rise to a direct relationship between illness outcomes and the match rate (i.e., regardless of vaccination).[18] For instance, if more probable strains are also more severe and vaccine strains are chosen based on expected severity, then the choice is more likely to generate a match when the upcoming influenza strain is more severe.[19] On the other hand, this process

[15] Regions are subprovincial and defined according to the geographic classification used by Statistics Canada. These "economic regions" are made up of adjacent census divisions meant to characterize regional economic activity. There are 73 economic regions in Canada, and after dropping the missing regions in Quebec and non-Winnipeg Manitoba, the total comes to 49.

[16] Specifically, the clinical match rate is calculated as the proportion of strains tested by the PHAC in each province and season that are matched (i.e., have cross immunity) with vaccine strains specified in that season:

$$m_{py} = \frac{number\ of\ flu\ tests\ in\ p,\ y\ that\ match\ vaccine\ strains\ selected\ in\ y}{number\ of\ flu\ tests\ in\ p,\ y}.$$

[17] Where labor force survey data is used, the unit of analysis is at the individual-region-week level, and additional controls include age, education, occupation, marital status, and union status.

[18] For more information on selection see: http://www.cdc.gov/flu/about/qa/vaccine-selection.htm (access Dec 21, 2012) or http://www.who.int/influenza/vaccines/virus/recommendations/en/index.html (access Dec 21, 2012).

[19] More specifically, suppose the WHO's choice of vaccine strains is driven by point-in-time forecasts of the future harm of circulating strains, where harm is informed by both the predicted severity of each strain and the perceived probability that each strain will appear. Consider the case where more probable strains are also more severe and the WHO chooses based on expected severity. Here, the choice more often leads to a good match when the upcoming influenza strain is more severe. This will drive a relationship between the match and the severity of the flu season regardless of vaccination (e.g., in the case where no one is vaccinated, a "would be match" is more likely to occur in a more severe flu season).

Table 1—Current Influenza Vaccination Status and the Clinical Vaccine Match Rate

	Overall		Interaction for subgroup			
			Age 65+	Poor self rated health	Chronic condition	Worker
Current match	−0.0015	−0.0010	0.0068	0.0001	0.0094	0.0118
	(0.0156)	(0.0167)	(0.0183)	(0.0184)	(0.0179)	(0.0312)
Lagged match		0.0015				
		(0.0155)				
Current match × subgroup			−0.0220	−0.0002	−0.0298	0.0102
			(0.0381)	(0.0378)	(0.0291)	(0.0387)

Notes: This table reports OLS estimates of the effect of the clinical match rate on current vaccination status. Each set of estimates shown in columns 1 to 6 represents a separate regression. Standard errors are given in parentheses and are clustered by economic region. Columns 3 to 6 include interactions of the current match with the subgroup listed in the table heading. Age 65+ is an indicator for age greater than or equal to 65. Poor self-rated health is an indicator of self-rated health equal to "poor" or "fair." Chronic condition indicates the presence of one of the following chronic conditions: Asthma, Heart Disease, High Blood Pressure, Diabetes, Cancer, or Emphysema/Chronic Bronchitis. Worker indicates that the respondent worked at a job or business in the last year. All regressions are weighted by survey weight and include province and year fixed effects.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

Source: Author calculations, PHAC surveillance reports, and the master files of NPHS cycle 2 and CCHS 1.1, 2.1, 3.1.

relies on accurate forecasting of strain likelihood and severity in early spring, and selection relies on known strains at that time. Meanwhile, new strain threats may occur while manufacturing is underway or complete. In short, it may be difficult to operate a policy to maximize the "true" cost of future strains, since information is restricted to speculative apriori probabilities of known strains. It is more likely that the observed empirical relationship between the clinical match rate and health operates through vaccination behavior. That is, when a matched strain appears, the vaccinated will enjoy a higher degree of protection from influenza relative to the unvaccinated, whereas when an unmatched strain appears, both the vaccinated and unvaccinated remain susceptible.[20]

Because the relationship between the match and strain severity is not directly testable, the empirical strategy is structured to explicitly account for this possibility. First, the model includes the vaccine match in level form to control for the direct effect of the vaccine match on illness. Second, the coefficient β_3 captures differences in the effect of the vaccine match post-2000 and hence controls for gains in the effect of the match that are common to all regions. Finally, the coefficient β_4 captures baseline differences in the effect of the match that are different in UIIC versus non-UIIC regions (i.e., regions may be differentially impacted by a higher

[20] While a direct effect of the match on the decision to vaccinate would have implications for this relationship, evidence indicates the match rate shows little predictive power for vaccination behavior. For instance, Table 1 presents results on the estimated impact of the match rate on vaccination status. Results indicate that there is little relationship between the current yearly match rate and current vaccination status, both overall and for the subgroups shown. In general, the estimated effect of a 1 percentage point increase in the match rate is associated with a statistically *insignificant* 0.0015 percentage point decrease in the likelihood of vaccination. Furthermore, last year's match rate shows little predictive power for current vaccination.

match even though vaccine composition each year is identical across all regions). Consequently, this methodology disentangles the relative change in the effect of the vaccine match that accrues to UIIC regions post-adoption. Specifically, β_1 summarizes the post-period difference in the effect of the match for UIIC regions, and captures the gain in health in good match years that is explained by increased vaccination from the program.

A. *Decreasing Returns*

The research design in equation (2) can be modified to capture decreasing returns to vaccination, which necessarily occur if externality benefits decline as vaccination levels rise. To address this aspect of vaccination programming, the model exploits variability in baseline average vaccination across regions. Since there is little between-variation in the post-program vaccination change for UIIC regions, post-program levels are heavily influenced by baseline vaccination differences (as shown in Table 2). To assess the role that baseline vaccination plays in the overall effect of the vaccination campaign, the model in (2) is modified to include interactions between baseline vaccination and all innovations of *UIIC*, *Post*, and *m*:

$$
\begin{aligned}
(3)\quad y_{jt} = {} & \beta_1[UIIC_p \times Post_y \times m_{py}] \\
& + \delta_1[UIIC_p \times Post_y \times m_{py} \times (preV_j - \overline{preV_j})] \\
& + \text{other interactions and controls} + \nu_{jt}.
\end{aligned}
$$

Here, $preV_j$ is the pre-2000 vaccination rate for region j, which is then demeaned by the overall average baseline vaccination rate. The coefficient β_1 is interpreted as the program match effect at the average baseline level, whereas δ_1 gives the change in the program match effect at higher baseline vaccination rates and estimates the degree to which decreasing returns impact program benefits. This strategy assumes that variability in baseline vaccination is related to post-program differences in health across good match and bad match years solely through decreasing returns to vaccination. Specifically, it assumes that the effect does not operate through unobservable characteristics among high and low regions that are otherwise related to future health outcomes. However, even if these unobservable factors are correlated with differences in baseline vaccination among regions, in order to explain results, such factors would also need to be related to future differences in health outcomes for UIIC regions, *and* differences in the way that UIIC regions persist through future good versus bad match flu seasons. Given that future vaccine matches are difficult to predict, that the vaccine is the same across North America, and that the program was largely unanticipated, the most likely explanation for UIIC regions with different baseline starting points to have different post-period changes in the effect of vaccine quality, is through decreasing returns to vaccination.

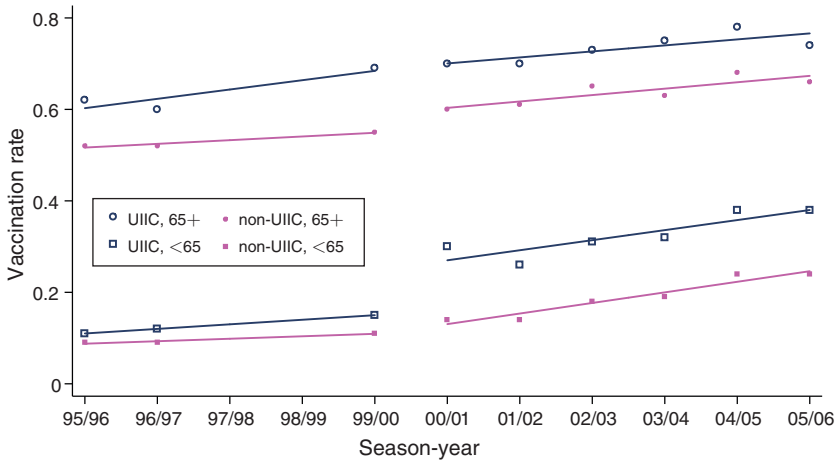

FIGURE 1. INFLUENZA VACCINATION—1995/1996 TO 2005/2006

Notes: The *y*-axis plots average vaccination for each season-year based on the master files of the National Population Health Survey Cycle 2 and the Canadian Community Health Survey Cycles 1.1, 2.1, 3.1. Solid lines show the fitted linear prediction for each age and geographic grouping.

III. Data and Descriptive Statistics

A. *Vaccination*

Health survey data from Statistics Canada are used to document the changes in vaccination. Four sequential health surveys contain questions relevant to influenza vaccination in the master file data holdings maintained by Statistics Canada's Research Data Centers. These are the National Population Health Survey (NPHS), Cycle 2 and the Canadian Community Health Survey (CCHS), Cycles 1.1, 2.1, and 3.1. Each survey is a national, population-based survey conducted on persons 12 years of age or older. Information includes demographic, socioeconomic, and health information, and current and previous vaccination status. In each survey, the respondent is asked: "Have you ever had a flu shot?" following up with "When was your last flu shot?" Along with survey timing, this information is used to determine the vaccination status of individuals in the current flu season-year (using the conventional definition of the season running from week 35 to week 34 of the following year).

Figure 1 shows vaccination rates for UIIC versus non-UIIC regions by age group over time. The figure shows that vaccination increased over 1996 to 2006 for all age groups, and indicates that the young have lower vaccination rates than the old over the sample period. The figure also shows that, while baseline vaccination rates for the young are not substantially different for UIIC versus non-UIIC regions, there are significant gains at program introduction as indicated by the sustained 11 percentage point relative shift in vaccination. The same is not evident in the older age group, where there is no relative change in vaccination at the introduction of the program.

TABLE 2— INFLUENZA VACCINATION RATES BY SELECTED CHARACTERISTIC

	UIIC			Non-UIIC			D-in-D	
	Pre	Post	Change	Pre	Post	Change	change	SD
Full sample								
Average vacc. rate	0.208	0.417	0.209	0.170	0.291	0.122	0.087	(0.023)
Between ER SD	0.021	0.032	0.017	0.018	0.031	0.019		
Average rate by age group								
Under 65	0.120	0.333	0.213	0.092	0.197	0.105	0.108	(0.023)
65 and over	0.609	0.737	0.128	0.521	0.645	0.124	0.004	(0.029)
Average rate by demographic characteristic and health status for under 65 group								
Male	0.111	0.292	0.181	0.076	0.169	0.093	0.088	(0.021)
Female	0.128	0.369	0.241	0.107	0.222	0.115	0.127	(0.023)
No secondary graduation	0.167	0.338	0.172	0.103	0.168	0.065	0.106	(0.025)
Secondary graduation	0.101	0.308	0.207	0.075	0.162	0.087	0.120	(0.019)
Some post-secondary	0.100	0.290	0.190	0.093	0.171	0.079	0.111	(0.021)
Post-secondary graduation	0.113	0.346	0.233	0.097	0.231	0.135	0.098	(0.024)
Income <$30K	0.130	0.350	0.220	0.101	0.186	0.085	0.135	(0.021)
Income $30K–$50K	0.118	0.336	0.218	0.087	0.192	0.105	0.113	(0.022)
Income >$50K	0.109	0.324	0.215	0.085	0.205	0.120	0.095	(0.024)
Full-time worker	0.096	0.301	0.205	0.079	0.187	0.108	0.097	(0.021)
Part-time worker	0.112	0.342	0.230	0.087	0.188	0.101	0.129	(0.021)
Not in labor force	0.187	0.448	0.262	0.149	0.263	0.114	0.147	(0.023)
No chronic conditions	0.085	0.267	0.182	0.061	0.143	0.082	0.100	(0.020)
At least one condition	0.237	0.479	0.241	0.201	0.326	0.126	0.116	(0.028)
SRH: excellent/very good	0.095	0.302	0.207	0.072	0.174	0.102	0.105	(0.023)
SRH: good	0.138	0.356	0.217	0.103	0.209	0.106	0.112	(0.021)
SRH: fair/poor	0.256	0.455	0.199	0.213	0.304	0.091	0.107	(0.024)
Full sample observations	40,012	119,294	159,306	31,824	144,774	176,598	335,904	

Notes: Vaccination rates for each subgroup are shown pre- and postprogram in UIIC and non-UIIC regions in Canada (excluding Quebec, rural Manitoba, and the Territories). Pre and post denote before and after September 2000 with the difference-in-differences change in vaccination for UIIC regions displayed in the second last column (standard errors clustered by economic region (ER) are shown to the right of the estimate). The between ER standard deviation calculates the standard deviation of vaccination rates between economics regions within each pre-post, UIIC-non-UIIC grouping. Chronic conditions include Asthma, Heart Disease, High Blood Pressure, Diabetes, Cancer, Emphysema/Chronic Bronchitis. SRH stands for self-rated health.

Source: Author calculations and the master files of NPHS cycle 2 and CCHS 1.1, 2.1, 3.1.

To explore differences in program impact among subgroups, Table 2 gives a summary of vaccination rates pre- and post-September 2000. The table shows an increase in vaccination rates for all regions in the post period with a 20.9 percentage point increase in UIIC regions and a 12.2 percentage point increase in non-UIIC regions. In relative terms, there is an 8.7 percentage point relative increase in UIIC regions, solely due to the increase of 10.8 for those under 65. Figure 2 confirms these results over the age distribution, showing that while UIIC regions have higher baseline levels of vaccination over all age groups (and in particular for those greater than 65), the relative change post-2000 is due solely to those under 65. It is clear from these data that the impact of the program is centered on the age group that was targeted by program incentives.

To better understand the characteristics associated with baseline vaccination behavior and heterogeneity in take-up after the program, the remainder of Table 2 presents summary statistics for ages under 65 by demographic group. As shown,

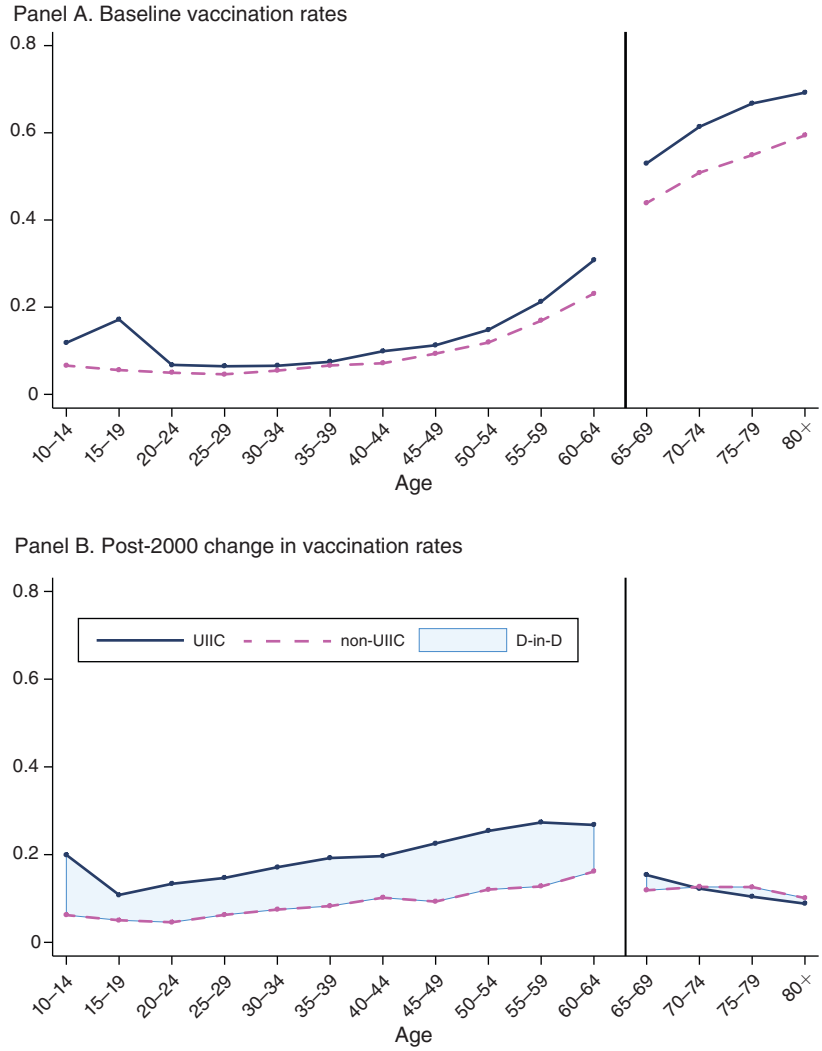

FIGURE 2. INFLUENZA VACCINATION RATES OVER AGE

Notes: In the top panel, the *y*-axis plots average vaccination for each age group pre-2000. The bottom panel shows the change in vaccination in the preperiod versus postperiod. Rates are based on data from the master files of the National Population Health Survey Cycle 2 and the Canadian Community Health Survey Cycles 1.1, 2.1, 3.1.

baseline vaccination patterns fall in line with previous research on the determinants of vaccination: females are more likely to vaccinate than males, vaccination increases with education (with the exception of those without secondary graduation), and decreases with time spent working (Mullahy 1999). Underlying health may play a part in explaining these patterns since health and socio-demographics are correlated, and the table shows that poor health is associated with higher vaccination rates. These baseline differences in vaccination will matter for overall program returns to the extent that immunity through vaccination varies by health status. For instance, if those in poor health have lower levels of vaccine efficacy then there may be room for the UIIC program to yield further external benefits for these groups.

Focusing on heterogeneity in vaccine take-up after the universal campaign, the table shows that among those under 65, the relative increase in vaccination for UIIC regions is largest among females, those not in the labor force, and among those with lower income. The relative increase is of similar magnitude regardless of having a covered chronic condition and across levels of self-rated health. It appears that the UIIC led to larger increases among more financially vulnerable groups, potentially ameliorating the gradient between SES and health. Overall, the most apparent source of heterogeneity is by age. Figure 2 confirms that, for those under 65, the highest degree of take-up is among the youngest and oldest groups.

B. *Influenza Surveillance*

Influenza surveillance data are used for two purposes in the analysis that follows. The first is to document the different influenza strains present in each season and province, and the second is to pinpoint the period in the year when influenza is circulating. In defining the epidemic period, results can be further delineated and used to verify that these benefits are delivered primarily when the threat of influenza is the greatest.[21] This offers a further specification check since, in order to explain results, any alternative factor would have to have a differential impact in the epidemic period *and* occur solely in UIIC regions, weeks post September 2000, and years with a higher vaccine match.

The surveillance data for laboratory confirmed influenza are available from the PHAC through its respiratory surveillance program. These data consist of weekly tests collected from appointed sentinel physicians in each surveillance region, where each test is then confirmed for presence of influenza or other respiratory diseases. These data are presented graphically in Figure 3. The bottom panel gives the average laboratory confirmed surveillance rate (percent of collected tests that are positive for influenza) and indicates the epidemic period of each year. Panel A of Figure 3 shows the clinical match rate for each season. To define the clinical match rate, I use strain isolation data from the PHAC along with reports on the cross-immunity of the yearly vaccine. All influenza strains observed by the PHAC are identified as matched or not matched to the yearly vaccine by using the reports published each year in the Canadian Communicable Disease Report (CCDR).[22]

Comparing panels A and B, the pattern between the vaccine match and influenza demonstrates higher peaks of laboratory confirmed influenza in seasons with a mismatch. Overall, across province and week, the estimated effect of the clinical match rate on the laboratory confirmed surveillance rate is -0.010 (s.e. 0.004), which is composed of an effect of -0.097 (s.e. 0.021) during the epidemic-period and -0.001 (s.e. 0.002) during the off epidemic-period.[23] The average clinical match rate is 0.71 and the average surveillance rate is 0.218 during the epidemic period

[21] The epidemic period is given as the contained set of weeks starting from the first week the number of positive influenza tests is greater than 5 percent of the season total until the last week it falls below 5 percent. Results are similar to defining the influenza period as any week with positive surveillance tests.

[22] These findings correspond to reports made by the CDC and the vaccine recipe from the WHO.

[23] These results employ the panel of surveillance data for week and province. Here, standard errors are clustered by province.

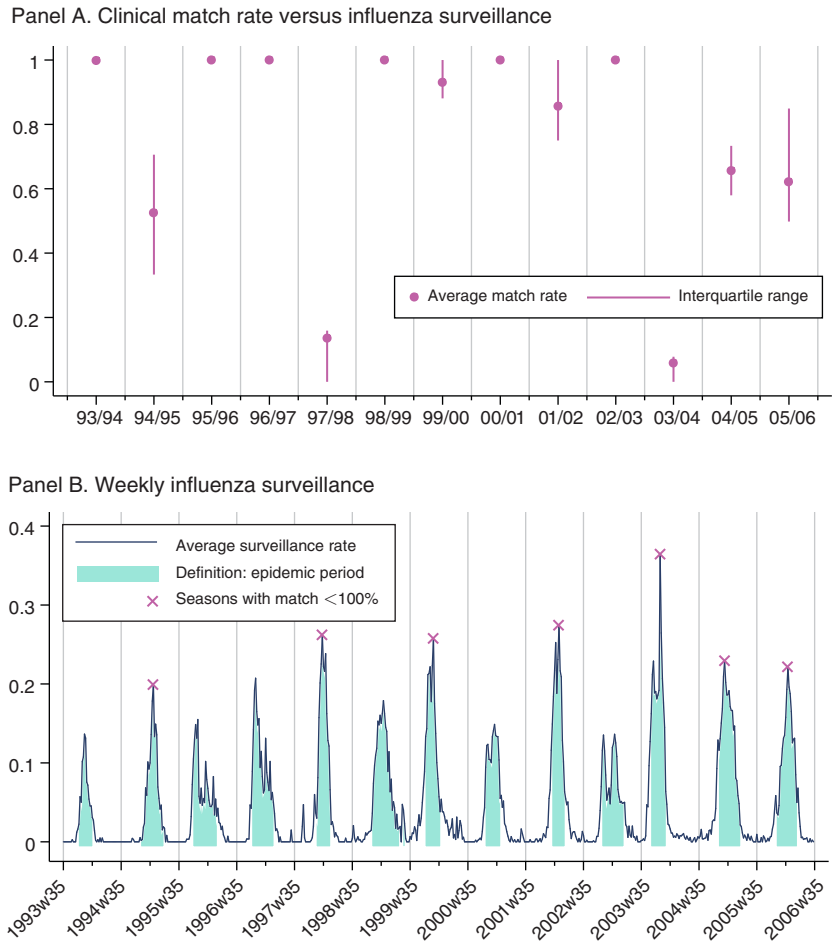

FIGURE 3. CLINICAL MATCH RATE AND INFECTIOUS DISEASE SURVEILLANCE

Notes: In panel A, the average clinical match rate is shown for each season with bars indicating the interquartile range (variation across province). The clinical match rate is calculated from the Canadian Communicable Disease Report and strain isolation data from the PHAC. Panel B shows the average fraction of positive influenza tests for data collected through the PHAC disease surveillance program. A tick at week 35 indicates the start of each season, the tick at week 9 indicates the season midpoint, and the shaded area indicates the epidemic period.

(0.024 during the off period). This implies that, during the epidemic period and relative to an average match, a perfect clinical match will result in a 2.8 percentage point decrease in the surveillance rate (or a 13 percent decrease from the mean).[24]

[24] Similarly, the match rate, itself, is directly related to the health outcomes studied below. For instance, the simple match effect (i.e., regression of outcomes on the match rate) is a decrease of 0.16 percentage points for work absences, 22.4 percent for influenza admissions and 12.7 percent for influenza-pneumonia admissions (all statistically significant at the 5 percent level).

C. *Health Outcomes*

Data on health outcomes is gathered from hospital administrative records and monthly labor force surveys. Hospital records are obtained from the Hospital Morbidity Database (HMDB) holdings of the Canadian Institute for Health Information. The HMDB includes complete records of inpatient discharges for hospitals in Canada.[25] Each discharge abstract consists of information on patient age, sex, and home postal code, as well as detailed medical information including: date of hospital admittance and discharge, discharge disposition (i.e., living or deceased), and detailed diagnosis information. Each abstract reports one diagnosis labeled the most responsible diagnosis (MRD) and up to 15 co-diagnoses.

The HMDB is used to construct weekly hospitalization rates for economic regions in Canada. Each economic region (ER) is a standard geographic region defined by Statistics Canada and made up of a group of adjacent census divisions meant to characterize regional economic activity. Focusing on ERs as opposed to census metropolitan areas (CMAs) has the advantage of covering the entire Canadian geography, both within and surrounding CMAs. Furthermore, it allows for control of fixed unobservable differences between well-defined regions, where region boundaries are set to capture localized economic activity (and likely tracking patterns of influenza transmission). Illness measures are assigned to each of 49 ERs based on patient postal code of residence.

Additionally, the monthly Labor Force Survey (LFS) is used to investigate effects for labor productivity. The LFS collects monthly information on demographic and labor market variables for household members 15 years of age and older. Demographic characteristics include age, sex, marital status, educational attainment, and family characteristics, and labor force characteristics including employment information, such as usual hours of work, and work absences in a reference week. Work absences are defined as "due to own illness."

Individual characteristics for work absences and hospitalizations are summarized in Table 3. The table shows that, among workers, those who report an absence in a given week are comparable to workers who do not report an absence (those reporting an absence are only slightly older and more likely to be female). Among hospital admissions (which, unlike the working population, represents the full age distribution), influenza admissions are more likely young and female compared to nonrespiratory admissions, while pneumonia admissions (a common complication of influenza and potential substitute diagnosis) are more likely older and male. Whereas influenza hospitalizations are shorter than nonrespiratory admissions, pneumonia admissions are of longer duration (likely reflecting respiratory infections of a more serious nature). Pneumonia admissions are also more likely to occur among residents of a residential care facility and end in death (the in-hospital death rate is 14.1 percent).

Weekly rates of illness are summarized in Table 4. On any given week, almost 5 percent of the labor force is absent from work for illness. The weekly rate of influenza is 0.4 per 100,000, but this average masks the degree of within season variation in this type of admission (within season variation is explored in what

[25] Hospitals in Quebec and non-Winnipeg Manitoba started submitting to the HMDB after 2001 and are consequently excluded from the analysis.

TABLE 3— SUMMARY STATISTICS–INDIVIDUAL CHARACTERISTICS

	Labor force absence		Hospital admission			
	Absence	No absence	Influenza	Pneumonia	All respiratory	Non-respiratory
Age	41.914	41.614	51.185	61.419	55.136	58.321
	(0.0015)	(0.0003)	(0.158)	(0.029)	(0.017)	(0.009)
Fraction male	0.494	0.594	0.440	0.535	0.529	0.518
	(0.0001)	(0.0000)	(0.002)	(0.001)	(0.000)	(0.000)
Duration	10.884		6.871	12.187	9.860	8.841
	(0.0012)		(0.083)	(0.024)	(0.012)	(0.007)
Wait time in ER			4.969	5.251	5.188	4.834
			(0.056)	(0.010)	(0.006)	(0.004)
Fraction with death			0.032	0.141	0.091	0.044
			(0.001)	(0.000)	(0.000)	(0.000)
Care home resident			0.058	0.099	0.063	0.038
			(0.001)	(0.000)	(0.000)	(0.000)
Urban postal code			0.775	0.862	0.856	0.868
			(0.002)	(0.000)	(0.000)	(0.000)
Observations	239,927	5,138,743	40,213	987,710	3,143,365	6,463,446

Notes: Variable means displayed with standard error of the mean given in parentheses below. Here, absence indicates a short-term work absence during the reference week. Nonrespiratory diagnoses include all hospitalizations that do not list a respiratory diagnosis as either an MRD (most responsible diagnosis) or as a contributing diagnosis. Duration is measured in hours for work absence and in days for hospital admissions. Wait time in ER (emergency room) before admission is measured in hours. Statistics from the LFS are weighted by survey weights.

Source: Author calculations, hospitalization data from the HMDB, and labor force data from the monthly LFS.

TABLE 4—COMPARISON OF ILLNESS POST-UIIC ACROSS HIGH AND LOW MATCH INFLUENZA SEASONS

	Overall illness rate	Post-UIIC effect			Difference in Post-UIIC effect
		Full sample	Low match	High match	High vs. low match
Work absence (× 100)	4.47	0.155**	0.388***	0.018	−0.370**
		(0.057)	(0.109)	(0.156)	(0.137)
ln(Influenza)	0.35	0.152*	0.243**	−0.092	−0.334***
		(0.079)	(0.104)	(0.066)	(0.091)
ln(FP)	8.62	−0.032	0.032	−0.078**	−0.110*
		(0.037)	(0.059)	(0.035)	(0.058)
ln(Respiratory)	26.18	−0.087	−0.042	−0.124**	−0.082
		(0.057)	(0.075)	(0.051)	(0.050)
ln(Nonrespiratory)	53.60	−0.083	−0.072	−0.068	0.004
		(0.054)	(0.071)	(0.042)	(0.041)

Notes: Differences in means are displayed for each variable with standard errors of each estimate given in parentheses below. All standard errors are clustered by region. The first column displays average illness rates for each variable (where HMDB admission rates are expressed as admissions counts per 100,000 population). Columns under the heading "Post-UIIC" display difference-in-differences estimates for each variable post September 2000 in UIIC regions. The final column displays the triple difference estimate of the Post-UIIC effect over seasons with a high match compared to low match. Work absence indicates a short-term work absence during the reference week, where estimates are multiplied by 100. Hospitalization data are expressed as the natural log of admission counts per week, per region. Statistics are weighted by survey weights in the case of the LFS, and population counts per year, per region in the case of the HMDB.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Source: Author calculations, hospitalization data from the HMDB, and labor force data from the monthly LFS.

follows). Expanding the diagnosis code to include potential substitute codes, the tables show average rates of 8.6 for influenza and pneumonia (FP) admissions and 26.2 for respiratory admissions. This compares to a rate of 53.6 for nonrespiratory admissions (these type of admission do not list a respiratory diagnosis any where on the discharge abstract).

To provide a preliminary understanding of the impact of the UIIC program, the table also breaks out the difference in illness for UIIC regions in the post period. In the absence of other differential factors in illness, this relative difference will reflect the true effect of the UIIC program. From Table 4 and with this interpretation in mind, we would conclude that the UIIC program led to *increases* in work absences and influenza admissions, while at the same time leading to *decreases* in both respiratory and nonrespiratory admissions, where the decreases in respiratory and nonrespiratory admissions are of comparable size, around 8 percent. What this interpretation ignores, however, is differential changes in other factors, such as labor conditions, the supply of hospital services, or diagnosis coding. In fact, when combining both influenza and pneumonia admissions there is a *decrease* in admits, which runs counter to the *increase* found for influenza alone and may partly reflect a higher likelihood of reporting a diagnosis of influenza versus pneumonia in post program UIIC regions.

To reconcile these results, the next two columns of Table 4 report post-UIIC estimates for high match seasons versus low match seasons, and the fourth column reports the difference in these estimates. The results indicate that high match seasons (i.e., a clinical match rate of one) are systematically associated with differential gains in respiratory health for UIIC regions in the post program period. For instance, overall there is an 8 percent relative decrease among all respiratory admissions in UIIC regions, but breaking this out over high and low match seasons shows that the decrease in admissions is primarily delivered in high match seasons. On the other hand, there is also an 8 percent relative decrease among *nonrespiratory* admissions in UIIC regions, but there is little difference in the decrease for seasons with a high clinical match versus those with a low clinical match.

IV. Results

A. *Main Results*

Summary evidence indicates that the clinical match rate is related to patterns of respiratory illness and suggests that this relationship operates through changes in vaccination programming. To present a more compelling case, the following will show the change in the effect of the vaccine match over different periods throughout the year. If a change in the effect of the match is related to vaccination and respiratory health, we would expect post-UIIC regions to have higher gains in high match years and, further, that these gains would be delivered in the typical epidemic period.

The results in Table 5 support both claims. First, over the entire year, UIIC regions have significantly lower illness rates in good match seasons following the vaccination program. Alternatively, non-UIIC regions have little difference in illness rates

TABLE 5—CHANGE IN POST-MATCH EFFECT ACROSS REGION AND MONTH

	Overall illness rate	Post-match effect		Difference in post-match effect
		UIIC regions	non-UIIC regions	
Work absence (× 100)				
All months	4.47	−0.355	0.014	−0.370
		(0.129)**	(0.064)	(0.137)**
Sep-Oct-Nov	3.76	−0.373	−0.192	−0.181
		(0.083)***	(0.070)**	(0.105)*
Dec-Jan-Feb	5.11	−0.573	0.101	−0.674
		(0.185)**	(0.146)	(0.227)***
Mar-Apr-May	4.97	−0.243	0.210	−0.452
		(0.174)	(0.190)	(0.250)*
Jun-Jul-Aug	4.03	−0.044	−0.052	0.008
		(0.182)	(0.128)	(0.213)
ln(Influenza)				
All months	0.35	−0.416	−0.081	−0.334
		(0.073)***	(0.061)	(0.091)***
Sep-Oct-Nov	0.19	−0.318	−0.726	0.408
		(0.087)**	(0.194)***	(0.209)*
Dec-Jan-Feb	0.79	−0.461	−0.161	−0.300
		(0.063)***	(0.125)	(0.137)**
Mar-Apr-May	0.29	−0.498	0.927	−1.424
		(0.152)**	(0.103)***	(0.176)***
Jun-Jul-Aug	0.08	0.227	0.075	0.152
		(0.045)***	(0.049)	(0.065)**
ln(Influenza + Pneumonia)				
All months	8.62	−0.083	0.026	−0.110
		(0.024)**	(0.054)	(0.058)*
Sep-Oct-Nov	7.71	−0.107	−0.115	0.008
		(0.034)**	(0.070)	(0.076)
Dec-Jan-Feb	11.10	−0.115	−0.020	−0.096
		(0.023)***	(0.049)	(0.054)*
Mar-Apr-May	9.01	−0.025	0.205	−0.229
		(0.034)	(0.066)***	(0.073)***
Jun-Jul-Aug	6.41	−0.029	0.091	−0.120
		(0.022)	(0.079)	(0.081)

Notes: Differences in means are displayed for the full season and for four periods running from the season start in September. Typical epidemic timing occurs around the season mid-point in February (see Figure 3). Standard errors are given in parentheses to the right of each estimate and are clustered by region. The first column displays average illness rates (where HMDB admission rates are expressed as admissions counts per 100,000 population). Columns under the heading "Post-match effect" display difference-in-differences estimates for high match-post September 2000 seasons (i.e., a high match indicates a clinical match of 1). The final column displays the triple difference estimate of the post-match effect over UIIC regions compared to non-UIIC regions. Work absence indicates a short-term work absence during the reference week (estimates are multiplied by 100). Hospitalization data are expressed as the natural log of admission counts per week, per region. Statistics are weighted by survey weights in the case of the LFS, and population counts per year, per region in the case of the HMDB.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Source: Author calculations, hospitalization data from the HMDB, and labor force data from the monthly LFS.

across good and bad match seasons in the post program period.[26] Second, these gains are delivered during the typical epidemic period (i.e., in the season running September to September, laboratory surveillance is typically elevated around the mid-season mark in February (see Figure 3), as are work absences and respiratory admissions (see Table 5)). Looking at work absences in UIIC regions, there is a post-program gain in the effect of the match of 0.37 percentage points in the first three months of the season, followed by a larger gain of 0.57 corresponding to the typical epidemic peak period. The size of the gain falls in magnitude and statistical significance in the March to May period, and, finally, there is a small and insignificant effect in the summer months. Meanwhile, non-UIIC regions experience an average 0.19 percentage point gain in the first three months of the season, but estimates throughout the rest of the year indicate that this represents a zero-sum displacement of work absences over the season (i.e., overall illness does not change, but the epidemic period appears later in the year). These same patterns are apparent for hospital admissions. Here, for UIIC regions, high match seasons deliver larger gains post program, and these gains are largest in the typical epidemic period. Alternatively, for non-UIIC regions, high match seasons do not deliver overall postprogram gains but tend to displace admissions from the early part of the season to later months.

Considered separately, the estimates displayed for UIIC and non-UIIC regions cannot disentangle the effect of vaccination from a change in the direct effect of the clinical match over time (e.g., the results cannot refute that unmatched strains are more "severe" in later years leading to even higher rates of illness in low match years in the post period). In order to isolate the effect of the vaccination program, the third set of results in Table 5 compares the post-match effect across UIIC and non-UIIC regions (here the vaccine is identical across region, but vaccination rates are not). Results indicate that post program gains in the effect of a high match are significantly higher for UIIC regions, and, further, these gains are largest in the mid-season period.[27]

While these figures show compelling patterns in the effect of vaccine quality over the year, health outcomes can be more flexibly modeled using equation (2), which fully exploits variability in the match rate across season and province, controls for other time-region varying factors in X, and is adjusted to capture differences in the timing of the flu epidemic year-to-year. Further, the estimation design can be modified to test the sensitivity of results to changes in compensatory behaviors by controlling for circulation of other infectious disease. Of interest is the coefficient on $UIIC \times Post \times M$, the program match effect, which gives the relative difference in health outcomes for higher match seasons in post program UIIC regions.

Estimates for equation (2) are shown in Table 6 and, overall, this evidence highlights the following two results. First, program-related health improvements are

[26] For instance, for work absences, the post-match estimate for UIIC regions is a statistically significant change of −0.36 percentage points, which means that in post 2000 a high match is associated with relatively lower levels of work absence. For non-UIIC regions, the estimate for work absences is an insignificant change of 0.01. For hospital admissions, UIIC regions experience a statistically significant change in influenza and combined influenza and pneumonia of −41.6 and −8.3 percent respectively, whereas non-UIIC regions experience an insignificant change of −8.1 and 2.6 percent respectively.

[27] The timing of the season does shift somewhat from year to year (see Figure 3), a feature that is explored explicitly below.

TABLE 6—THE FLU IMMUNIZATION CAMPAIGN, VACCINE MATCH, AND HEALTH OUTCOMES

Dependent variables	Average duration (1)	Work absence (× 100) (2)	ln(influenza) (3)	ln(influenza + pneumonia) (4)
Panel A. Base results				
Epidemic-period	12 weeks	−0.522**	−0.949***	−0.486***
		(0.21)	(0.26)	(0.07)
Baseline rate		4.56	2.96	16.05
Off-period	40 weeks	−0.103	0.021	−0.056
		(0.38)	(0.11)	(0.05)
Baseline rate		3.55	0.19	7.69
Panel B. Account for behavioral response to match by controlling for circulation of other infectious disease				
Epidemic-period		−0.629***	−0.914***	−0.476***
		(0.24)	(0.26)	(0.07)
Off-period		−0.136	0.020	−0.054
		(0.11)	(0.11)	(0.05)

Notes: This table reports estimates of the interaction of the clinical match rate, a dummy for post September 2000, and a dummy for UIIC regions. Table columns report results for different health outcomes and rows report results for two different periods in the year: the epidemic-period and the off-period. Column 1 shows the average duration of each period. Work absence indicates a short-term work absence during the reference week (estimates are multiplied by 100). Hospitalization data are expressed as the natural log of admission counts per week, per region. Each estimate shown in columns 2 to 4 is from a separate regression. Standard errors are given in parentheses and are clustered by region. In panel A, baseline illness rates (i.e., rates for preprogram UIIC) are reported below the standard errors (where HMDB admission rates are expressed as admission counts per 100,000 population). Statistics are weighted by survey weights in the case of the LFS, and population counts per year, per region in the case of the HMDB. All regressions include the level effect of the match rate, month, age, season and economic-region fixed effects as well as interactions of *UIIC*×*Post*, *Post*×*Match*, and *UIIC*×*Match*. Regressions in columns 3 and 4 also control for public health expenditures on health care (hospitals, capital investments, physicians and other health professionals), and diagnosis specific coding classifications changes (ICD10 versus ICD9). Regressions in column 1 control for the same factors, and for education, marital status, sex, occupation and union status. Results in panel B control for surveillance rates of other infectious respiratory disease (respiratory syncytial virus, parainfluenza, and adenovirus).

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Source: Author calculations, hospitalization data from the HMDB, and labor force data from the monthly LFS.

delivered primarily in the epidemic period.[28] For instance, higher matches in post program UIIC regions deliver statistically significant decreases in work absences and respiratory admissions in the epidemic period, whereas there is little difference in all illness outcomes in the off-epidemic period.[29] Second, there is little difference in point estimates (in magnitude and significance) after controlling for other infectious respiratory disease (i.e., respiratory syncytial virus, parainfluenza,

[28] This pattern in results is not by design. Unless vaccination is a mitigating factor, there are few reasons why program-match estimates for absences and admissions should systematically follow the surveillance patterns of laboratory influenza counts.

[29] Since the log specification drops weeks without admissions, the results for influenza admission in the off-epidemic period do not include the full set of weeks. However, an alternative specification using influenza admission rates on the left-hand side and including all 0's shows similar results. There is a small and insignificant effect in the off-epidemic period of −0.006 per 100,000. In the case of combined influenza and pneumonia admissions, there are no weeks with zero admissions.

and adenovirus).[30] If compensatory behaviors (related to the match rate and post program UIIC) impact influenza circulation, then this should also relate to circulation of other infectious disease. However, controlling explicitly for circulation of these diseases leaves point estimates largely unchanged.

The magnitude of results points to large program effects during the epidemic-period. For instance, for work absences, the coefficient on $UIIC \times Post \times M$ is -0.6 percentage points (a 14 percent decrease from the baseline mean or six less absences per 1,000 workers, per week).[31] The implications for decreases in underlying influenza infections are most readily seen from the results on influenza admissions, where there is a 91 percent decrease arising from a high match under the UIIC program. Unless these results reflect large changes in the probability of hospitalization conditional on infection or changes in in-hospital diagnosis coding for influenza, they indicate near elimination of influenza infections at the population level in post program-high match seasons.

The next column expands the chosen diagnosis code to include pneumonia (a common complication of influenza and a potential substitute diagnosis). This both captures any underlying influenza-associated pneumonia diagnoses and, by comparing the magnitude of results, can shed light on whether the large effects results for influenza admissions are a feature of changes in diagnosis coding. Overall, estimates show a 48 percent decrease in FP admissions (implying a savings of 8.6 admissions per 100,000 per week or a fall of 2.5 admissions going from the average match, 71 percent, to a perfect match).

In order to compare these estimates to those for influenza admissions and to get a sense of what they imply for infection at the population level, the FP results need to be scaled to account for fact that not all FP admissions are due to influenza infection. While the proportion of infection-associated FP admissions is unknown, it can be approximated by comparing FP admits across off-epidemic and epidemic periods. For instance, off-epidemic FP admission rates are half that of epidemic rates at baseline. Similarly, the size of the program match effect in the epidemic period is a decrease of almost half. These results indicate that an excellent match under the UIIC program leads to near elimination of influenza associated FP admissions. Explicitly, using the relative difference in average FP rates reported in Table 6 as a scale factor, these results suggest a 92 percent decrease in influenza infection in the underlying population (comparable to the effect found for influenza admissions).[32]

[30] These diseases are nonvaccine preventable respiratory viruses that are infectious through the same manner as influenza and present with similar symptoms.

[31] Similar results arise for other health outcomes in the NPHS and CCHS health survey data. For instance, during the epidemic period, the program match effect for cold/flu medications is a decrease of 10.2 percentage points or a 13 percent decrease from the mean (interestingly, there is no impact on the use of antibiotics). Secondly, over the full sample, there is a 3.0 percentage point fall in the rate of being in bed ill. In contrast, the effect for workers is 0.6 percentage points, which is very similar to the effect found for absences in the LFS survey results in Table 6.

[32] This statement assumes, first, that the probability of hospital admission conditional on illness (either influenza or other) is not sensitive to the program match effect and, second, that the comparison between epidemic and off-epidemic rates is a good approximation of the relative rate of influenza infections for a given set of diagnosis codes. In support of these assumptions, evidence on *nonrespiratory* admits indicates little sensitivity to program match effect, and the "rates approximation" is likely not unreasonable given that the set of FP codes are still narrowly focused around influenza and the epidemic period is defined by positive influenza surveillance. Under

TABLE 7— THE INFLUENZA IMMUNIZATION CAMPAIGN, VACCINE MATCH, AND TIME LOST TO ILLNESS

Dependent variables	Work absence (1)	ln(influenza) (2)	ln(influenza +pneumonia) (3)
Panel A. Total time lost (*hours absent per respondent/total hospital days per region-week*)			
Epidemic-period	−0.046**	−0.901**	−0.349***
	(0.02)	(0.42)	(0.12)
Baseline rate	0.38	18.26	171.92
Off-period	−0.013	0.044	−0.005
	(0.03)	(0.21)	(0.06)
Baseline rate	0.26	1.66	98.52
Panel B. Time lost per illness (*hours absent per absence/hospital days per admission*)			
Epidemic-period	0.240	−0.066	0.067*
	(0.64)	(0.07)	(0.03)
Baseline rate	11.66	7.11	10.71
Off-period	−0.447	−0.058	0.033
	(0.32)	(0.10)	(0.02)
Baseline rate	11.28	6.94	12.21
Panel C. Deaths per region-week			
Epidemic-period			−0.386***
			(0.12)
Baseline rate			1.91
Off-period			−0.016
			(0.06)
Baseline rate			1.13

Notes: See Table 6 notes. Table rows report results for time lost due to illness for each of the heath outcomes listed in the table headings. The duration of work absences is given in hours and is collected from the LFS. In this case, panel A presents results for hours absent over the full sample, whereas panel B presents the results for hours absent conditional on reporting a work absence. The duration of hospital admissions is given in days and is collected from the HMDB. In this case, panel A presents results for the log number of hospital days in each region-week (with average rates given in hospital days per 100,000), and panel B presents results for the log number of hospital days per admission (with average rates given in hospital days per admission). Finally, panel C presents results for the log number of deaths in each region-week (with average rates given in deaths per 100,000). There are too few influenza deaths to present results for log influenza deaths.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

Source: Author calculations, hospitalization data from the HMDB, and labor force data from the monthly LFS.

Table 7 shows that the program match effect on work absences is mainly due to decreased absences rather than due to shorter hours spent ill per illness (in fact, hours spent per absence are slightly elevated in post-UIIC-match). Furthermore, the decrease in the number of hospital days is mainly due to fewer hospitalizations since average hospital length of stay increased only marginally (e.g., length of stay increased by 7 percent for FP, whereas overall hospital days fell by 35 percent). Lastly, there are large effects for death from FP; a high match after the immunization

these assumptions, we can benchmark the effect for infection in the population. Here, the percent change in the influenza infection rate arising from the program match effect will be equivalent to the percent change in hospital admission scaled by the inverse ratio of the probability of admission with infection to the probability of admission. For FP diagnoses, under the rates approximation, this is equivalent to $-0.48 \times 16.05/(16.05 - 7.69) = -0.92$.

TABLE 8— THE INFLUENZA IMMUNIZATION CAMPAIGN, VACCINE MATCH, AND RESPIRATORY VERSUS NON-RESPIRATORY ADMISSIONS

log number of admissions by type of diagnosis	MRD respiratory disease	Codiagnosis or MRD respiratory disease	No diagnosis of respiratory disease	All diagnoses
Epidemic-period	−0.366***	−0.261***	0.038	−0.088***
	(0.05)	(0.04)	(0.02)	(0.02)
Baseline rate	23.80	36.73	56.86	93.59
Off-period	−0.044	−0.033	0.029	0.010
	(0.04)	(0.04)	(0.02)	(0.02)
Baseline rate	14.49	27.08	57.03	84.10

Notes: See Table 6 notes. Hospital admissions are divided by whether a respiratory diagnosis is listed on the hospital abstract record (i.e., as "most responsible" or as a codiagnosis). Results for each category are presented using the log number of admissions in each region-week (with average rates reported below standard errors and given in admissions per 100,000).

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

Source: Author calculations and hospitalization data from the HMDB.

program reduced deaths by 39 percent. This represents a gain of almost 9 fewer deaths per 100,000 over the epidemic period.

B. *Impact on Respiratory versus Nonrespiratory Admissions*

Influenza infection may lead to further health complications for those in poor health or with preexisting conditions (e.g., chronic respiratory disease, heart disease, cancer, or nervous system diseases). To analyze the impact of the vaccine program on related co-morbidities, Table 8 shows results for all admissions where a respiratory diagnosis is listed as either the most responsible cause (MRD) or as an additional co-diagnosis. This has the advantage of not only capturing MRD respiratory diagnoses but also capturing MRD diagnoses for other conditions that are present with related respiratory complications. For comparison, the table also shows results for admissions where no related respiratory diagnosis is listed.

In the epidemic period, respiratory complaints comprise a large proportion of overall admissions. For instance, considering the baseline rates in Table 8, almost 40 percent of total admissions include respiratory as a contributory or MRD diagnosis (25 percent when considering only MRD respiratory). Comparing average admission rates across the epidemic and off-epidemic periods, the seasonal aspects of respiratory disease also become apparent. In the off-epidemic period MRD respiratory admits are 39.0 percent lower, and co respiratory/MRD admits are 26.3 percent lower. Meanwhile, nonrespiratory admits differ by only 0.3 percent in the epidemic versus off-epidemic period. In total, hospital admissions are 10.1 percent lower in the off-epidemic period.

Evidence points to large effects of the UIIC campaign in the epidemic period. The post-UIIC-match effect estimates a 36.6 percent reduction in MRD respiratory admissions and a 26.1 percent reduction in co-respiratory/MRD admissions (essentially equivalent to the difference in admission rates between the epidemic

and off-epidemic periods). Meanwhile, the program appears to have little effect on respiratory admits in the off-epidemic period, and little effect for nonrespiratory admits in any period. These results provide further support that findings are not explained by alternative factors related to health outcomes, since such factors (i.e., correlated with the vaccine match and specific to the time and place of the immunization program) would likely manifest in measures of health across the board. Conversely, the evidence shown here indicates that the impact of the program is specific to respiratory complications and moreover, follows the timing of the seasonal epidemic.[33]

Considering total admissions, the estimated program match effect is an 8.8 percent decrease. Considering off-epidemic admissions rates as a baseline for "non-epidemic" rates, the size of this estimate implies that the program combined with a high match largely eliminates influenza-associated admissions and, assuming no change in the probability admission conditional on illness (i.e., no crowd-out or crowd-in effect), suggests a 87 percent reduction in the probability of influenza infection at the population level.[34] Adjusting for crowd-in among nonrespiratory admission yields a benchmark of 95 percent (similar to the benchmark found using the results for FP admissions).[35]

C. *Impact by Age and Externality Effects for Older Adults*

The UIIC program led to a change in vaccine coverage for those under age 65 and, as shown in Figure 4, increases in vaccination post program center on this age group. While difference-in-differences estimates for age groups 65+ show little increase in vaccination (even slight decreases past age 70), the relative increases in vaccination among younger groups are on the order of 10 percentage points.

The bottom panel of Figure 4 compares changes in vaccination with the relative percent decrease in FP admissions in post UIIC high match years.[36] The largest decreases are among the very young (under ten); high school and college age; and prime age individuals age 45–55. Notably, however, there are sustained decreases in FP admissions in the 65+ group even though there is no matching increase in vaccination. Assuming this age group is unaffected by the vaccination of others, we would expect program benefits for ages 65+ to be close to zero.[37] However,

[33] The evidence on nonrespiratory admits also shows that there is little evidence of significant "crowd in" at the extensive margin among other disease admissions.

[34] This benchmark is calculated in the same manner as the benchmark found for FP admissions:
$-0.088 \times (93.59/(93.59 - 84.10)) = 87$ percent.

[35] The "crowd in" effect for nonrespiratory admits is small and statistically indistinguishable from zero, but since the point estimates are positive, adjusting for changes in the probability of admission given a nonrespiratory illness would only strengthen the evidence for large decreases in influenza. This is because the percent change in overall admissions would need to be adjusted downward to account for size of the crowd in effect. For instance, the relative increase in nonrespiratory admissions in the epidemic period is 0.9 percent and the ratio noninfluenza to influenza admissions can be approximated using the baseline rates in Table 8: $(56.89 + 27.08)/(36.73 - 27.08)$. The leads to a scaled adjustment of 8 percent, or, in other words, the estimated reduction in influenza infections would be $87 + 8 = 95$ percent.

[36] Since point estimates for the off-epidemic period are small and statistically insignificant, only estimates for the epidemic period are displayed in Figure 4.

[37] For instance, under this "no externalities" assumption we would expect the relative change in admissions to match the relative change in vaccination (the latter of which is zero for older groups). An alternative argument

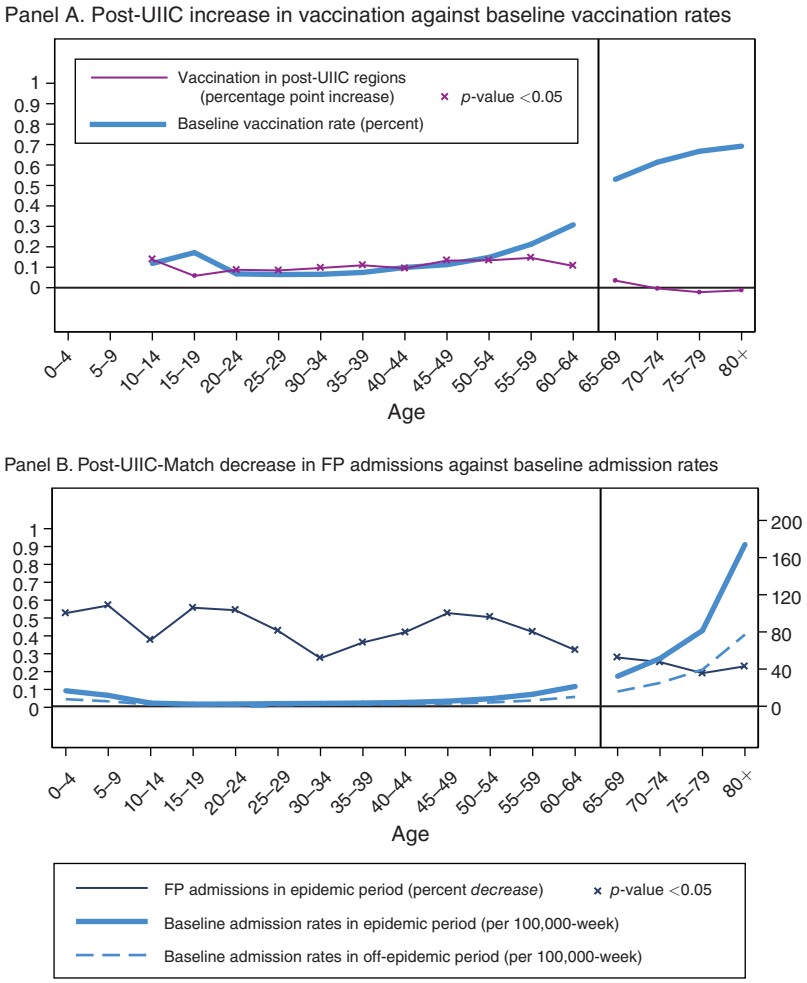

FIGURE 4. POST-UIIC CHANGE IN VACCINATION AND PROGRAM MATCH EFFECT FOR FP ADMISSIONS

Notes: The top figure plots the relative change in vaccination rates postprogram in UIIC regions (i.e., difference-in-differences estimates) by age group (dark line). For comparison, baseline vaccination rates (i.e., preprogram UIIC rates) are also given (light line). The bottom figure (left axis) plots the age specific coefficients on the program match effect (the interaction of the clinical match rate, a post-September 2000 dummy and a dummy for UIIC regions) for influenza and pneumonia admissions during the epidemic period (given as percent decrease). For comparison, the light blue line (right axis) shows baseline FP admission rates (per 100,000 per week).

the evidence indicates that those 65+ benefit from changes in vaccination among younger groups; even where relative vaccination levels for older age groups have marginally decreased, there are still large program benefits. Furthermore, these external benefits appear in spite of higher baseline vaccination levels among UIIC

is that since baseline rates of vaccination among older groups are already quite high, this leaves little scope for improvement among those older than 65 (both in relative and absolute terms). In either case, the data refute both hypotheses. Furthermore, evidence indicates that effectiveness of the vaccine among older groups is low (the leading explanation being poorer immune response among older groups; see, for instance Fiore et al. 2010 or Jefferson et al. 2007), which leaves further scope for improvement, by external means, even at high baseline levels.

regions (i.e., under decreasing returns to vaccination this would suggest *negative* relative benefits).[38]

D. *Benchmarking Averted Influenza Infections*

Extending this analysis to calculate total benefits in terms of averted infections requires further information about how the probability of hospital admission relates to influenza infection in the population. Unfortunately, such information is unobserved in the data. However, as before, we can benchmark this effect under the following two conditions: the rate of admission conditional on illness (influenza or otherwise) is not sensitive to the program match effect, and the difference between baseline epidemic and off-epidemic rates is a good approximation of the proportion of influenza-associated admissions.[39]

Using these conditions and starting with the estimates for all age groups, we can compare the inferred decrease in influenza infection to the change in the vaccinated population. First, note that if there is no difference in the probability of infection for remaining unvaccinated individuals (i.e., no externality or selection effects), then the overall probability of infection should fall only because more unvaccinated individuals are now vaccinated. In this case, (and assuming a matched vaccine) the percent change in the infection rate should equal, not exceed, the percent change in the size of the unvaccinated group.[40] However, evidence for the case of a matched vaccine suggests a decrease in the infection rate of 92 percent, which is much larger that the 11 percent decrease in the size of the unvaccinated group.[41]

[38] For instance, if there is decreasing returns to vaccination, then, given baseline levels, separate groups may experience similar increases in vaccination but expect different gains in health. Since baseline vaccination levels for older adults are higher in UIIC regions, then, under decreasing returns, this should lead to smaller relative gains. On the other hand, estimates point to larger relative gains for older adults in UIIC regions, indicating substantial external benefits from increased vaccination in younger groups. The role for decreasing returns to vaccination is explored further in the next section.

[39] The direction of the program match effect for admissions conditional on illness (i.e., the crowd in effect) is likely zero or positive, in which case the benchmark estimate will be a lower bound on the size of the effect. In fact, the program match effect for nonrespiratory indicates a small and statistically insignificant change in admissions rates. Extrapolating from here, it is likely that there is little or no change in the rate of "crowd in" for the more narrowly defined set of FP codes. In terms of the proportion of influenza-associated FP admissions, the rates approximation is likely reasonable given that the set of FP codes is still largely focused around influenza and that the epidemic period is defined by positive influenza surveillance.

[40] To see this, consider a simplified framework where the influenza infection rate can be written as a weighted average of the conditional infection rate for vaccinated and unvaccinated groups. If the vaccine is perfectly protective of the vaccinated group, then the infection rate is equivalent to the infection rate among the unvaccinated group multiplied by relative group size (this can be relaxed to the case of an imperfectly protective vaccine by modifying the size of the unvaccinated group by the degree of under-protection in the population and scaling estimates down accordingly). Take a change in the size of the unvaccinated group. If there is no difference in the probability of infection for remaining unvaccinated individuals, then the overall probability of infection will fall only because more unvaccinated individuals are now vaccinated (i.e., a composition effect). In this case, the percent change in the infection rate and the percent change in the size of the unvaccinated group should be comparable.

[41] Since this considers the case of a matched vaccine, the benchmark estimate of the decrease in infection is based on the UIIC program effects for FP admissions under a clinical match rate of one (see Section IVA). Meanwhile, the estimate used for the decrease in the unvaccinated is calculated using the 9 percentage point decrease in the unvaccinated relative to the baseline of 79 percent in preprogram UIIC (see Table 2). This yields an 11 percent decrease in the size of the unvaccinated population. This stylized framework scales the difference in the protection rate between vaccinated and unvaccinated groups by the size of the match rate (given as the clinical match in the empirical work) and abstracts from any other differences in protection rates among vaccinated and unvaccinated groups (e.g., imperfect immune response among the vaccinated or partial immunity among the unvaccinated).

Since the decrease in the unvaccinated population is not large enough to explain the large decreases in infection, this points to changes in infection rates for those with no change in vaccination behavior (either through externalities or through selection, the latter of which may arise when program-compliers have different baseline infection probabilities relative to non-compliers). Extant externality or selection effects are likely positive since a decrease in the size of the unvaccinated group likely decreases the spread of disease (externality effect), and those with higher chances of infection are likely the first to vaccinate (selection effect). Furthermore, the two effects need not be independent. Positive selection into vaccination could give rise to an "add on" externality effect if those with higher baseline rates of infection remove themselves first and therefore prevent the spread of disease at a higher rate.

Empirical evidence indicates the presence of one or both of these effects since the degree of averted infections is above and beyond what can be explained by the change in the size of the vaccinated population. Figure 5 provides further insight into these differences by making the same comparison over the age distribution. To see the extent of inferred excess benefits, the figure re-plots the change in FP admissions, but adds the relative decrease in the unvaccinated group (calculated from the estimates in Figure 4, top panel) and the benchmarked decrease in influenza infections (calculated from the FP estimates and relative baseline rates in Figure 4, bottom panel). Here, for those over 65, it is possible to isolate the externality effect since there is no relative difference in vaccination and, thus, little scope for selection. The comparison shows external benefits for older adults benchmarked at about 40 to 50 percent (i.e., comparing averted infections over decreases in the unvaccinated). Choosing an even more conservative scale factor, the externality effects are still apparent. For instance, under the assumption that 0 percent of FP admissions are due to other diseases (scale factor of one), the results show external benefits at a minimum of 20–30 percent.

Are large external benefits to older adults realistic in the face of already high vaccination rates? One interpretation of these results is that vaccination of the young yields higher returns when compared to the "own" effects of vaccination for this group. This accords with the medical literature, which indicates that effectiveness of the vaccine among older groups is low due to poorer immune response (Jefferson et al. 2007; Fiore et al. 2010) and provides support for policies aimed at vaccinating close contacts of care residents (Poland, Tosh, and Jacobson 2005).

Looking over the full age distribution, there are positive excess gains from the UIIC for all age groups. While the evidence for older adults indicates that externalities have a role to play, the results for those under 65 are a combination of both externalities and selection. Selection is a virtue of the UIIC anywhere that the change in behavior is more likely among those with higher relative rates of infection. Not only will this type of selection mechanically reduce the infection rate among those who remain unvaccinated, but it also leads to stronger externality effects if selection

If there is differential protection among vaccinated and unvaccinated groups, this could lead to additional selection effects, the implications of which are discussed below.

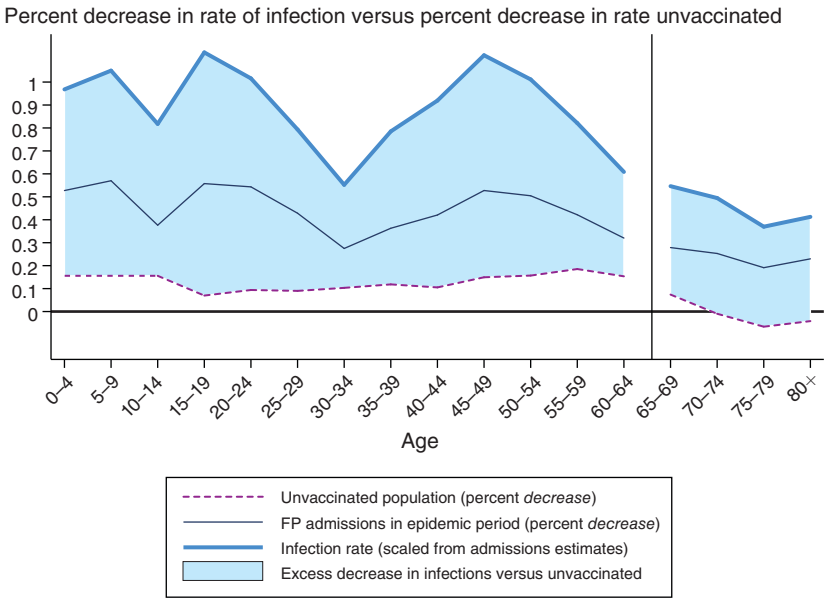

FIGURE 5. PROGRAM MATCH EFFECT AND INFERRED EXCESS BENEFITS

Notes: The bottom plot line gives the percent decrease in the unvaccinated population (the ratio of difference-in-differences estimates to baseline rates for the unvaccinated (see Figure 4 top panel)). The middle line replots the age specific coefficients on the program match effect (the interaction of the clinical match rate, a post-September 2000 dummy and a dummy for UIIC regions) for influenza and pneumonia admissions during the epidemic period (given as percent decrease). The top line plots the inferred decrease in influenza infections scaled from the estimates for FP admissions (the ratio of FP estimates to the proportional difference between baseline epidemic and off-epidemic rates (see Figure 4 bottom panel)). The health surveys did not collect information on vaccination for those 11 and under.

removes those who spread the disease at a higher rate.[42] Finally, baseline levels of vaccination will inform the overall size of excess gains to the extent that there are decreasing returns to scale. In this case, the large estimated magnitude indicates that the change in vaccination in UIIC regions occurred at a key baseline starting point, a feature explored explicitly below.

E. *The UIIC and Decreasing Returns*

Externalities have direct implications for the impact of vaccination if external spillovers eventually drive returns to zero. For instance, evidence based on approximations of the transmission rate and duration rate of influenza indicates that an infected individual mixing in a wholly unvaccinated population, would, on average, infect 1.44

[42] In light of the large externality effects found for those over 65, the externality effects among the young are likely nonzero. Moreover, the evidence provides a strong case against pure selection. For example, a pure selection effect of this size would follow if the change in vaccination occurred only among a small group of unvaccinated susceptibles (e.g., if only 10 percent are susceptible and only susceptibles choose to vaccinate, this would decrease the unprotected population by $0.09/0.10 = 90$ percent). However, to the extent that individuals have imperfect information about their own susceptibility, it is unlikely that this behavior fully explains results. Furthermore, Table 2 shows limited differences in the relative change in vaccination by health status, which suggests, perhaps, a more limited scope for pure selection effects.

others before recovery (Hethcote 2000; Boulier, Datta, and Goldfarb 2007). Using this "contact rate" within a standard model of disease dynamics implies that a vaccination rate greater than 31 percent would reduce the average infection number to below one, (assuming a fully protective vaccine).[43] Since an average infection rate below one will prevent an influenza epidemic (i.e., an infection less than replaces itself), it suggests that benefits from immunization beyond 31 percent would fall to zero.

The model described in Section IIA explores this possibility by allowing the program match effect to depend on baseline immunization levels. Results from this analysis are presented in Table 9, where panel A represents base results and panel B shows results allowing for the interaction of baseline immunization among regions. The results show that there are decreasing returns depending on baseline immunization starting points. Specifically, the results for $UIIC \times Post2000 \times Match$ show the effect at average baseline levels across region, whereas the results for the baseline interaction term indicate a decline in the impact of $UIIC \times Post2000 \times Match$ at higher baseline rates. The baseline rates where the program match effect would be zero are shown at the bottom of the table.[44] The numbers indicate that a region starting with vaccination rates around 32–35 percent, but receiving the same UIIC boost in vaccination, would enjoy little additional gain from a high match vaccine post program. On the other hand, based on the average starting point of 18 percent, the expected effect of the program is large relative to baseline incidence rates, indicating that the change in vaccination occurred at a highly effective point in the baseline distribution.

F. Estimated Benefits versus Program Costs

These results suggest that vaccination leads to substantial benefits in terms of hospitalization and lost work costs. For instance, estimates indicate that there is a 26 percent decrease in comorbidity/MRD respiratory hospitalizations, which implies a decrease of 9.6 respiratory hospitalizations per 100,000 per week during the epidemic period. Over the population of Ontario, this is a savings of 1,245 admissions per week. The length of the average epidemic period is 12 weeks and the cost of an average respiratory hospitalization is $8,629 (CIHI 2008). A back of the envelope calculation indicates a savings of $129 million when the match is high. Expected savings at the average match rate are $92 million. For illness absences, estimates indicate a 0.6 percentage point decrease in work absences. Over the working population of Ontario, this is a savings of 47,400 work absences per week during the epidemic period. At an average hourly wage of $18 and average absence duration of 10.9 hours this translates into $112 million in savings per season for Ontario in high match years (expected savings of $79 million). These

[43] This model is based on the Kermack and McKendrick Susceptible-Infective-Removed (SIR) model of disease epidemics. Several variations of the model are shown in Geoffard and Philipson (1997); Francis (1997, 2004); and Boulier, Datta, and Goldfarb (2007). Using the contact number of 1.44, the dynamics of this model imply that the average number of infections occurring from one infection will fall to one when at most 69 percent of the population is susceptible to influenza (i.e., $1.44 \times 0.69 = 1$). In other words, to prevent an epidemic, the SIR model implies that more than 31 percent of the population would need to be vaccinated (with a high match vaccine).

[44] Using the parameters in equation (3), these rates are calculated by setting $\beta_1 + \delta_1 \times (preV - 0.182)$ equal to zero. This implies that the program match effect is zero when baseline rates equal: $-\beta_1/\delta_1 + 0.182$.

Table 9—Results Conditional on Baseline Vaccination Rates in Each Economic Region

Dependent variables	Work absence (× 100) (1)	ln(influenza) (2)	ln(influenza +pneumonia) (3)
Panel A. Results for epidemic weeks			
UIIC × *Post* × *Match*	−0.629***	−0.914***	−0.476***
	(0.24)	(0.26)	(0.07)
Panel B. Results for epidemic weeks conditional on baseline vaccination in each economic region			
UIIC × *Post* × *Match*	−0.712***	−0.816**	−0.511***
	(0.10)	(0.38)	(0.06)
UIIC × *Post* × *Match* × *Baseline vaccination*†	5.145	4.978	3.544**
	(3.69)	(7.58)	(1.68)
Baseline rate where program effect is zero	32.0%	34.6%	32.6%

Notes: See Table 6 notes. Panel A reports estimates of *UIIC* × *Post* × *Match* from equation (2) for all weeks during influenza season. Panel B reports estimates of *UIIC* × *Post* × *Match* and an estimate of *UIIC* × *Post* × *Match* interacted with the demeaned baseline vaccination rate as specified in equation (3). In panel B, the baseline vaccination rate displayed is the implied baseline rate at which the effect of the influenza program would be zero, given a perfect vaccine match. It is calculated as the baseline average vaccination rate plus the negative ratio of the two estimates above the result. Table columns report results for different health outcomes. Work absence indicates a short-term work absence during the reference week (estimates are multiplied by 100). Hospitalization data are expressed as the natural log of admission counts per week, per region. Standard errors are given in parentheses and are clustered by region. Statistics are weighted by survey weights in the case of the LFS, and population counts per year, per region in the case of the HMDB.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

†Baseline vaccination denotes the deviation of the rate for each region from the average baseline rate of 0.182.

Source: Author calculations, hospitalization data from the HMDB, and labor force data from the monthly LFS.

savings are less than programming costs, which are an average $19 million in additional administration costs and an extra $14 million in vaccine costs for a total additional cost of $33 million per season.

V. Conclusion

This paper illustrates the role of externality effects in immunization behavior. Principally, where externalities exist, the limitations of previous evidence to address these externalities manifests along two dimensions: unaccounted spillovers in estimation and failure to recognize systematic differences in overall benefits depending on the level and composition of local immunization. Evidence that overlooks these aspects can lead health authorities to set immunization policies well before full benefits are recouped, and, furthermore, to neglect the implications of policies set within the context of differing underlying immunization levels. For instance, previous literature indicates mixed evidence on the benefit of vaccinating younger adults, but largely ignores estimation spillovers and is silent on how benefits depend on local immunization levels and compositional take-up.

This paper addresses these aspects by studying a novel immunization program expanding vaccination coverage outside the standard target group. Within this context, evidence demonstrates significant overall and external gains of immunization, and shows that these gains decline to exhaustion depending on local

immunization levels. First, overall benefits under UIIC are substantial, centering primarily on children and older adults. Second, there is direct evidence of external benefits for those over 65 arising from increased vaccination of younger groups. In terms of the UIIC, the gains in vaccination occurred at the most effective range possible. Overall, vaccination rose from an average of 20.9 percent to 41.4 percent, and evidence suggests that gains from further increases in vaccination would be small.

Compositional effects in vaccine take-up are material to the interpretation of these results, and where selection exists, the implications of evidence acquired from randomized design may be limited, particularly where policy cannot easily replicate random assignment. Importantly, the nature of the UIIC program generates significant differences in take-up across age groups. This is partly by design, given previous immunization policy for older groups, but is also apparent for those under 65. Selection of this nature has direct implications for overall expected benefits. For instance, holding the local vaccination rate constant, overall benefits may differ under scenarios where different groups are targeted or selection differs within group. Specifically, immunization policies, informed by previous evidence, focused on maximizing the immunization of those over the age of 65. The results presented here show that, even with immunization rates for older adults at levels in excess of 60 percent, there is still room for substantial gains. Recognizing the role for externalities and selection in overall immunization benefits may shed light on the effectiveness of this particular policy target. In particular, evidence indicates that the UIIC program is an effective tool to deliver benefits to older adults, demonstrating that large gains accrue from: heterogeneity in the effect of immunization for younger groups, combined with further amplified external benefits delivered by younger groups. Importantly, since differences in cost effectiveness calculations between the young and the old may reflect, primarily, differences in treatment cost upon infection (which run opposite the immune protection afforded by the vaccine), ignoring external benefits delivered from young to old can severely undervalue the overall cost effectiveness of vaccination of the young.

Since the evidence presented demonstrates the important role of externalities in the expected impact of immunization programs, it suggests that future research might benefit from a more explicit approach to assessing the implications of these aspects, both in terms of estimation (i.e., addressing treatment spillovers) and in overall policy design (i.e., understanding how benefits depend on the local immunization rate and compositional take-up). In particular, the UIIC program is one that expanded recommendations and coverage outside the typical target group. Using this evidence, firms or governments may wish to reconsider these recommendations and the potential gains from providing coverage outside this target group. Furthermore, evidence presented is shown to depend on compositional take-up among younger groups, and external benefits delivered within the context of the local immunization rate. Within this policy context, results indicate that program expansion leads to substantial decreases in illness, which aside from direct and indirect benefits to individuals, delivers large relative cost savings to the publicly financed health care system. Lastly, since results imply that high vaccination rates combined with a matched vaccine can lead to large gains, this may suggest a further advantage to developing seasonal infrastructure that can be relied upon in the case of a matched pandemic vaccine (for example, the 2009 H1N1 vaccine).

REFERENCES

Almond, Douglas. 2006. "Is the 1918 Influenza Pandemic Over? Long-Term Effects of *In Utero* Influenza Exposure in the Post-1940 U.S. Population." *Journal of Political Economy* 114 (4): 672–712.

Almond, Douglas, and Bhashkar Mazumder. 2005. "The 1918 Influenza Pandemic and Subsequent Health Outcomes: An Analysis of SIPP Data." *American Economic Review* 95 (2): 258–62.

Benson, Veronica, and Marie A. Marano. 1998. *Vital and Health Statistics: Current Estimates From the National Health Interview Survey, 1995*. U.S. Department of Health and Human Services. Series 10 (199) Data from the National Health Survey. Washington, DC, October.

Boulier, Bryan L., Tejwant S. Datta, and Robert S. Goldfarb. 2007. "Vaccination externalities." *B. E. Journal of Economic Analysis & Policy* 7 (1): Article 12.

Canadian Institute for Health Information. 1996–2006. "Hospital Morbidity Database." Canadian Institute for Health Information. http://www.cihi.ca/CIHI-ext-portal/internet/en/document/types+of+care/hospital+care/acute+care/HMDB_METADATA.

Centers for Disease Control and Prevention. 2011. *Prevention and Control of Influenza with Vaccines: Recommendations of the Advisory Committee on Immunization Practices* (*ACIP*), *2011*. Centers for Disease Control and Prevention. August.

Centers for Disease Control and Prevention. 2012. *Prevention and Control of Influenza with Vaccines: Recommendations of the Advisory Committee on Immunization Practices* (*ACIP*)*—United States, 2012–13 Influenza Season*. Centers for Disease Control and Prevention.

Demicheli, Vittorio. 2001. "Mass influenza vaccination in Ontario: Is it worthwhile? " *Canadian Medical Association Journal* 164 (1): 38–39.

Fiore, Anthony E., David K. Shay, Karen Broder, John K. Iskander, Timothy M. Uyeki, Gina Mootrey, Joseph S. Bresee, and Nancy J. Cox. 2008. *Prevention and Control of Influenza with Vaccines: Recommendations of the Advisory Committee on Immunization Practices* (*ACIP*)*, 2008*. Centers for Disease Control and Prevention. Atlanta, August.

Fiore, Anthony E., David K. Shay, Penina Haber, John K. Iskander, Timothy M. Uyeki, Gina Mootrey, Joseph S. Bresee, and Nancy J. Cox. 2007. *Prevention and Control of Influenza: Recommendations of the Advisory Committee on Immunization Practices* (*ACIP*)*, 2007*. Centers for Disease Control and Prevention. Atlanta, July.

Fiore, Anthony E., Timothy M. Uyeki, Karen Broder, Lyn Finelli, Gary L. Euler, James A. Singleton, John K. Iskander, et al. 2009. *Prevention and Control of Influenza with Vaccines: Recommendations of the Advisory Committee on Immunization Practices* (*ACIP*)*, 2009*. Centers for Disease Control and Prevention. Atlanta, August.

Fiore, Anthony E., Timothy M. Uyeki, Karen Broder, Lyn Finelli, Gary L. Euler, James A. Singleton, John K. Iskander, et al. 2010. *Prevention and control of influenza with vaccines: recommendations of the Advisory Committee on Immunization Practices* (*ACIP*)*, 2010*. Centers for Disease Control and Prevention. Atlanta, August.

Francis, Peter J. 1997. "Dynamic epidemiology and the market for vaccinations." *Journal of Public Economics* 63 (3): 383–406.

Francis, Peter J. 2004. "Optimal tax/subsidy combinations for the flu season." *Journal of Economic Dynamics and Control* 28 (10): 2037–54.

Geoffard, Pierre-Yves, and Tomas Philipson. 1997. "Disease Eradication: Private versus Public Vaccination." *American Economic Review* 87 (1): 222–30.

Govaert, Th. M. E., C. T. M. C. N. Thijs, N. Masurel, M. J. W. Sprenger, G. J. Dinant, and J. A. Knottnerus. 1994. "The Efficacy of Influenza Vaccination in Elderly Individuals." *Journal of the American Medical Association* 272 (21): 1661–65.

Groll, Dianne L., and David J. Thomson. 2006. "Incidence of influenza in Ontario following the Universal Influenza Immunization Campaign." *Vaccine* 24 (24): 5245–50.

Health Canada. 2007. *Access To The Seasonal Flu Vaccine In Canada: How the flu shot makes its way from the laboratory to the doctor's office.* Health Canada. Ottawa.

Hethcote, Herbert W. 2000. "The Mathematics of Infectious Diseases." *SIAM Review* 42 (4): 599–653.

Jefferson, Tom, Carlo Di Pietrantonj, Alessandro Rivetti, Ghada A. Bawazeer, Lubna A. Al-Ansary, and Eliana Ferroni. 2010. "Vaccines for preventing influenza in healthy adults." Cochrane Database of Systematic Reviews (Online) July 7(CD001269).

Johansen, Helen, Kathy Nguyen, Luling Mao, Richard Marcoux, Ru-Nie Gao, and Cyril Nair. 2004. "Influenza vaccination." *Health Reports* 15 (2): 33–43.

Kermack, W. O., and A. G. McKendrick. 1927. "A Contribution to the Mathematical Theory of Epidemics." *Proceedings of the Royal Society of London, Series A: Mathematical, Physical, and Engineering Sciences* 115 (772): 700–21.

Kurji, Karim. 2004. "The Ontario Experience with Universal Vaccination." Paper presented at National Influenza Vaccine Summit, Atlanta, April.

Kwong, Jeffrey C. 2005. "The effect of Ontario's Universal Influenza Immunization Program on pneumonia and influenza hospitalizations." MS diss. University of Toronto.

Kwong, Jeffrey C., and Douglas G. Manuel. 2007. "Using OHIP physician billing claims to ascertain individual influenza vaccination status." *Vaccine* 25 (7): 1270–74.

Kwong, Jeffrey C., Thérèse A. Stukel, Jenny Lim, Allison J. McGeer, Ross E. G. Upshur, Helen Johansen, Christie Sambell, William W. Thompson, Deva Thiruchelvam, Fawziah Marra, Lawrence W. Svenson, and Douglas G. Manuel. 2008. "The Effect of Universal Influenza Immunization on Mortality and Health Care Use." *PLoS Med* 5 (10): e211.

Manzoli, Lamberto, Francesco Schioppa, Antonio Boccia, and Paolo Villari. 2007. "The Efficacy of Influenza Vaccine for Healthy Children: A Meta-Analysis Evaluating Potential Sources of Variation in Efficacy Estimates Including Study Quality." *Pediatric Infectious Disease Journal* 26 (2): 97–106.

Miguel, Edward, and Michael Kremer. 2004. "Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities." *Econometrica* 72 (1): 159–217.

Ministry of Health and Long Term Care. 2000. "Ontario invests 38 million to ease emergency room pressures with universal vaccination program." *Canada News Wire*, July 25. http://ogov.newswire.ca.ontario/GPOE/2000/07/25/c6018.html?/1match=&lang=_e.html.

Mullahy, John. 1999. "It'll only hurt a second? Microeconomic determinants of who gets flu shots." *Health Economics* 8 (1): 9–24.

National Advisory Committee on Immunization (**NACI**)**.** 1995–2006. *Statement On Influenza Vaccination For The Season*. Canadian Communicable Disease Report Vol. 21–32. Ottawa.

National Advisory Committee on Immunization (**NACI**)**.** 2012. *Statement on Seasonal Influenza Vaccine for 2012–2013*. Canadian Communicable Disease Report 38 (ACS-2). Ottawa, August.

National Center for Health Statistics, and National Health Interview Survey. 2008. Public-use data file and documentation. http://www.cdc.gov/NCHS/nhis/nhis_2008_data_release.htm, 2008.

Nichol, Kristin L. 2001. "Cost-Benefit Analysis of a Strategy to Vaccinate Healthy Working Adults Against Influenza." *Archives of Internal Medicine* 161 (5): 749–59.

Nichol, Kristin L. 2003. "The efficacy, effectiveness and cost-effectiveness of inactivated influenza virus vaccines." *Vaccine* 21 (16): 1769–75.

Nichol, Kristin L. 2008. "Efficacy and effectiveness of influenza vaccination." *Vaccine* 26 (Supplement 4): D17–22.

Philipson, Tomas. 2000. "Economic Epidemiology and Infectious Diseases." In *Handbook of Health Economics,* Vol. 1B, edited by Anthony J. Culyer and Joseph P. Newhouse, 1761–99. Amsterdam: Elsevier.

Poland, Gregory A., Pritish Tosh, and Robert M. Jacobson. 2005. "Requiring influenza vaccination for health care workers: seven truths we must accept." *Vaccine* 23 (17–18): 2251–55.

Public Health Agency of Canada (**PHAC**)**.** 1995–2006. *Influenza Season: Canadian Laboratory Diagnoses And Strain Characterization.* Canadian Communicable Disease Report Vol. 21–32. Ottawa.

Public Health Agency of Canada (**PHAC**)**.** 2006. "Canadian Immunization Guide." http://www.phac-aspc.gc.ca/publicat/cig-gci/index-eng.php.

Statistics Canada. 1995–1999. "National Population Health Survey (NPHS), Cycle 2: master files." Statistics Canada Health Statistics Division. Ottawa.

Statistics Canada. 1996–2006. "Labour Force Survey master files." Statistics Canada Labour Statistics Division. Ottawa.

Statistics Canada. 2000–2006. "Canadian Community Health Survey (CCHS) Cycle 1.1, 2.1, 3.1: master files." Statistics Canada Health Statistics Division. Ottawa.

World Health Organization. 1996–2006. "Mortality Database." World Health Organization (WHO). Geneva, CH. http://apps.who.int/healthinfo/statistics/mortality/whodpms/.

World Health Organization. 2003. *Influenza: Report by the Secretariat to the 56th World Health Assembly.* World Health Organization (WHO) A56/23. Geneva, CH, March.

Thompson, William W., David K. Shay, Eric Weintraub, Lynette Brammer, Nancy Cox, Larry J. Anderson, and Keiji Fukada. 2003. "Mortality Associated With Influenza and Respiratory Syncytial Virus in the United States." *Journal of the American Medical Association* 289 (2): 179–86.

Thompson, William W., David K. Shay, Eric Weintraub, Lynette Brammer, Nancy Cox, Larry J. Anderson, and Keiji Fukada. 2004. "Influenza-Associated Hospitalizations in the United States." *Journal of the American Medical Association* 292 (11): 1333–40.

Ward, Courtney J. 2014. "Influenza Vaccination Campaigns: Is an Ounce of Prevention Worth a Pound of Cure?: Dataset." *American Economic Journal: Applied Economics.* http://dx.doi.org/10.1257/app.6.1.38.

Zhang, Gang, Dany Shoham, David Gilichinsky, Sergei Davydov, John D. Castello, and Scott O. Rogers. 2006. "Evidence of Influenza A Virus RNA in Siberian Lake Ice." *Journal of Virology* 80 (24): 12229–35.