NBER WORKING PAPER SERIES

# LEARNING BY SUFFERING? PATTERNS IN FLU SHOT TAKE-UP

Ginger Zhe Jin
Thomas G. Koch

Working Paper 25272
http://www.nber.org/papers/w25272

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2018

We thank Lenisa V Chang and participants at the 2016 ASHEcon conference for helpful comments, and the Federal Trade Commission for research support. All errors are ours. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Learning by Suffering? Patterns in Flu Shot Take-up
Ginger Zhe Jin and Thomas G. Koch
NBER Working Paper No. 25272
November 2018
JEL No. D12,D18,I12,I18

## ABSTRACT

An annual flu shot is one of the least controversial and most widely-recommended preventative health measures. In spite of such advice, only a fraction of those who are suggested to get a flu shot actually receive it. We focus on past personal outcomes to understand how individual learning in influences patterns over time using medical claims for a 5% panel sample of Medicare FFS beneficiaries. We find that individuals learn from personal suffering from flu and such learning is conditional on whether they have taken a flu shot in the same flu season. If they did not take a flu shot, having the flu later on encourages them to get the flu shot next year. But if they had the flu shot and still got the flu, their likelihood of getting a flu shot next year is significantly reduced. The 2009 outbreak of bird flu does not break the qualitative pattern of "learning by suffering" but it does change the strength of learning.

Ginger Zhe Jin
University of Maryland
Department of Economics
3115F Tydings Hall
College Park, MD 20742-7211
and NBER
jin@econ.umd.edu

Thomas G. Koch
FTC
600 Pennsylvania Avenue, NW
Washington DC 20580
tgkoch@gmail.com

# 1 Introduction

An annual flu vaccine (or shot) is one of the least controversial and most widely-recommended preventative health measures. The Center for Disease Control's Advisory Committee on Immunization Practices (ACIP) suggests that everyone six months or older get an annual flu vaccine, absent pre-existing allergies or severe health conditions. This advice is consistent with the vaccine's well-established effectiveness at limiting both private and social risks of getting the flu.

Despite the strength of these guidelines, the adherence rate for this advice is imperfect, and, as we document below, likely less than previously thought.(Fiscella et al., 2006) Further, using panel data, we demonstrate a striking pattern of vaccination over time. In particular, we find that individuals in our sample "learn from suffering." The learning is both positive and negative: when individuals are un-vaccinated and subsequently get sick, they are more likely to get a flu vaccine the following year. Alternatively, when individuals are vaccinated and get the flu in spite of the vaccine, they are less likely to be vaccinated the following flu season. This "learning by suffering" is robust across a series of specifications, even after we control for individual fixed effects plus the region's previous flu shot rate and flu illness rate.

The 2009 outbreak of bird flu does not break the qualitative pattern of "learning by suffering" but it does change the strength of learning. This change suggests that individuals interpret the signal of flu illness differently after they observe a new strain of flu virus, but the new interpretation does not last long. It is most prominent in 2010, the first flu season after the bird flu outbreak. As the threat of bird flu subdued, the strength of "learning by suffering" gradually recovers towards its original levels before 2009.

To our best knowledge, we are the first to document "learning by suffering" of vaccination using individual-level claims data. While "learning by suffering" is intuitive, it is also surprising for this population: since the recommendation of flu vaccination has been widespread for a long time, one would imagine that substantial learning, if ever needed, has already taken place for most adults before age 65. It is also surprising that the learning is contingent on individual experience, even after we control for region-wide factors. One explanation is that self experience is most salient to individuals, individuals are uncertain

about the strength of their own immune system, and the effectiveness of flu shot is perceived to vary across individuals despite the universal recommendation of flu shot. These individualized uncertainties encourage people to learn from self experience rather than take the blanket recommendation at its face value.

According to the US Center of Disease Control, 66.2% of the elderly (65 or older) received flu vaccine in the 2012-2013 flu season.[1] In contrast, the average vaccination rate recorded in our 5% sample of Medicare claim data is less than 40% every flu season from 2005-2006 to 2012-2013. While this discrepancy could be driven by some elderly getting flu vaccine without filing a claim, it is also possible that some subjects over report vaccination when they were asked whether they had received any flu shot in the past 12 months. Such over-reporting, if it exists, may lead CDC to over-estimate the actual vaccination rate.[2] While this potential over-reporting is of interest to us, our main focus is on dynamic patterns of learning that are a source of under-vaccination.

We also shed light on the potential mechanisms involved in individuals making flu shot decisions. Moyer (2018) argues that the recent anti-vaccination movement has made it difficult to publish findings that highlight the limitations of vaccination. That is, the scientific community fears individuals might over-respond to credible scientific research that, while still valuable, a flu shot is not as effective in the general population as previously thought. Medicare beneficiaries respond strongly to their own experience in the immediate past, despite their life-time exposure to the CDC recommendation.

The rest of the paper is organized as follows. Section 2 articulates our contribution to the literature. Section 3 describes the data and related institutions. Section 4 presents the empirical estimates. Section 5 discusses the implication of our findings for public health policy.

[1]CDC report accessed at http://www.cdc.gov/flu/fluvaxview/coverage-1213estimates.htm on June 1, 2016.

[2]The CDC estimates are based on the National Immunization Survey conducted by the CDC National Center for Immunization and Respiratory Diseases (NCIRD) and the Behavioral Risk Factor Surveillance System (BRFSS) survey conducted by various states.

# 2 Literature Review

Most of the economic literature of vaccine focuses on externality. Because vaccines protect against infectious diseases, the vaccinated population generates a positive externality on non-vaccinated people, while a non-vaccinated person may contract the disease and impose a negative externality on others. Assuming vaccination incurs more private cost (in time, discomfort, and side effects) than the marginal, private benefit of vaccination, the last unvaccinated individual has strong incentive to free ride on the vaccinated population. As a result, the whole population is likely under-vaccinated absent government intervention.

In response, scholars have sought incentive-changing policies to mitigate under-provision. Most economic literature of vaccine focuses on a positive externality from vaccinated to non-vaccinated population. (Chang, 2015; Ward, 2014; Chapman et al., 2012; Brenzel et al., 2007; Bronchetti et al., 2015) Some even advocated for tort remedy, which holds non-vaccinated people liable when they transmit a vaccine-preventable disease to others.(Reiss, 2013) The efficacy of policy-driven incentives depends on the efficacy of vaccines, which may vary both objectively and subjectively across individuals. Recent research has addressed how these two factors interact and agree that individuals should receive external incentives to get vaccinated, though the magnitude of such incentives is a point of contention.(Naprawa and Reiss, 2014; Manski, 2010, 2014; Gostin, 2015; Kahan et al., 2010; Ryan and Nourmohammadi, 2010)

In contrast, our paper focuses on individual incentives while controlling for geographic variation. More specifically, each state and county in the US estimate their region demand for flu shot each year and guide flu shot supply in advance. Beside nationwide coordination by the US Center of Disease Control, local governments also engage in local educational campaigns in each flu season, some even provide free flu shots to the public via community events. It is unclear whether these efforts are motivated by the free-riding problem, or simply by an aim to reduce the cost of flu shot to individuals. Either way, we can control for a combination of individual fixed effects and year fixed effects, and region-specific history of flu and flu shot.

Our study is closer to the literature that strives to explain low vaccine take up at the

individual level. Some research has focused on two potential mechanisms to resolve the problem: providing information and reminders to patients (Milkman et al., 2011), and healthcare provider communication.(Opel et al., 2015)[3]

Past immunization against influenza can be predicted by subject-estimated effectiveness of the vaccination, periodic blood test, perceived severity of flu illness, side effects of vaccine, health anxieties, and subjective probability of being infected.(Tsutsui et al., 2008) Others have focused on the importance of opportunity cost, as captured by labor-force participation. (Carman and Mosca, 2011) Risk perception has also been used to understand the take-up rates of HPV vaccines.(Kahan et al., 2010)

There is evidence that inattention is not a leading cause of under-vaccination, though that work is focused on a different, younger population than we consider. (Bronchetti et al., 2015) Rather, informed decisions to not get the vaccine, and lack of follow-through are more important. The regulatory environment(McConeghy, 2014; Calandrillo, 2003) may also play a role.

To our knowledge, we are the first to empirically document a novel mechanism in vaccination take up, namely the *evolution* of individual beliefs on their own disease susceptibility and vaccine effectiveness. Our research design allows us to identify the effect of learning while controlling for chronic health condition, risk preference, and opportunity cost via individual fixed effects. In addition, we explore how the learning process changes before and after the 2009 outbreak of bird flu, which complements an experimental study on belief coordination after the outbreak. (Engle-Warnick et al., 2013) Our region-wide controls also capture aggregate peer effects inside a particular region, no matter whether individual vaccination decisions are negatively correlated due to free-riding or positively correlated due to the desire to conform to peers.(Bodine-Baron et al., 2013; Rao et al., 2007)

Absent a thorough understanding of individual behavior, attempts to fix it via public health awareness campaigns or more effective vaccines may not work or be counterproductive. A better understanding of individual take-up behavior also helps on the supply side. (Finkelstein, 2004; Scherer et al., 2007; Dai et al., 2015) In complement, our findings

[3]There is a broader literature on encouraging preventive care, with or without externalities, cf. Stone et al. (2002).

suggest that demand for vaccines is related to individual perception of vaccine quality and the evolution of individual perception has lasting effects in future demand of vaccines.

# 3 Data and Institutional Detail

We start with the baseline population, which is a 5% random sample of Medicare fee-for-service (FFS) population for the years 2005-12.[4] We follow this people from when they enter Medicare FFS until they exit it, primarily due to death. We restrict our study sample to those individuals who are 65 or older, and whose primary eligibility criteria for Medicare FFS is old age. Medicare FFS also covers individuals on SSI disability and end-stage renal dialysis. These individuals enter Medicare FFS due to a pre-existing health problem, and are excluded. For the sample, we observe each beneficiary's age, gender, race and ZIP code of residence.

We observe billed medical events for our 5% sample. Principally, we consider claims for services offered in three main settings: hospital inpatient admissions[5], and outpatient care offered in either hospital or non-hospital settings, such as ambulatory surgical centers or a doctor's office.[6] All claims list two key pieces of information: CPT codes describing the services provided, and ICD-9 codes characterizing the health circumstances of the individual. The claims list the date(s) of service, provider identifiers (NPI or UPIN), and biller type—doctor's office, clinic, or, importantly, bulk billers, typically used for mass immunization.

We use these CPT and ICD-9 codes to identify the utilization of flu shots, as well as incidence of the flu. The various codes for flu shots and the flu are distinct; that is, the codes for treatment for flu-like symptoms are different than those for administration of the flu shot. Medicare billing practice has established several key codes in effect for billing these services during this time period. These codes are listed in Table 1. The bird flu outbreak of 2009-10 is included in our sample; CMS specified several new ICD-9 codes to reflect this

[4]One notable consequence of this is that we do not observe the medical claims for individuals who enter into the Medicare Part C, or Medicare Advantage, program.

[5]Technically, we use the MedPAR files, which are summary files for all inpatient admissions and skilled-nursing facility stays.

[6]We also have access to the prescription drug use by way of Medicare Part D plans starting in 2006. Flu shots are billed under Medicare Part B; if a patient received a prescription for, e.g., Tamiflu, to treat flu-like symptoms, the visit where the prescription was written would be in the claims data.

and other new forms of the flu. We also include the broader category of respiratory illness, which may be confused by patients or practitioners for the flu.

We construct flu seasons that begin in August and end the following July. With the exception of the 2009-10 bird flu season, this is consistent with the CDC's construction of flu seasons and reflects patterns of flu-shot take-up and flu evident in the data. Figure 1 plots the number of claims for flu shots and flu-related (by ICD-9 code) claims. Nearly 80 percent of flu shots are deployed by November, and flu outbreaks tend to begin in January. This likely reflects the two-week lag between when the flu shot is administered and when it can credibly provide protection against the flu. That is, someone who receives the flu shot in the midst of an outbreak will not be protected against that outbreak for some period of time. This is why the CDC emphasizes timely flu shots early in the season. Because of this timing, we will include flu shot outcomes in the 2012 season, since most of the flu shots likely have been taken by the end of our data (December 31, 2012). We do not use the outcomes on the flu using the 2012 flu season measures, since it is much more likely to be incomplete.

The one interesting exception to this pattern was the 2009 bird flu season. As seen in Figure 2, the seasons before and after 2009 exhibit the standard timing (i.e., flu-related medical visits peak some time after peak vaccine administration), while the bird flu season outbreaks began around the time that the flu vaccine was being distributed. Further, this irregular timing contributed to the poor match between the flu strains in standard vaccine and those that were prevalent in the outbreaks. Below, we will evaluate if the take-up and other behavioral patterns varied pre- and post-2009 season.

Sample means are reported in Table 2. The take-up rate for our sample (roughly 35-40 percent by year) is substantially lower than self-reported measures for a similar population (age 65+). There are two potential reasons for this gap. The first and more concerning one is that individuals may get flu shots outside of the Medicare billing system. Seniors may choose to pay out of pocket to a provider, even though they do not face a co-pay for flu shots in Medicare FFS. (The provider might refuse Medicare FFS, and expect more money out-of-pocket from a beneficiary than from CMS.) County health systems may be one source, though CMS, through its bulk billing system, does provide reimbursement services for these kinds of providers. Employers might also offer flu shots outside of the billing system, but

individuals of this age are not likely to be employed.

Alternatively, it may be the case that self-reported measures of preventative care overstate the actual utilization. This appears to be true for a wide array of preventative health care services, such as mammograms, Pap smears and cholesterol tests, when claims-based measures are compared to self-reported measures in the Medicare population for services that are covered and there is little reason to leave the Medicare insurance billing system (Fiscella et al., 2006).

Similar issues may arise for our measure of flu illness. We code an individual as having had the flu if one of her medical claims during a flu season has ICD-9 codes for the flu. A person may experience flu-like symptoms, not be formally diagnosed with the flu, but believe they suffered from the flu and behave accordingly. Alternatively, a beneficiary may be diagnosed with the flu on the medical claim, but not be aware of it, or not believe the diagnosis when told by the provider.

Principally, this is important because we are measuring the response of the individual to realizations (or perceptions thereof) of the risk of the flu that they face, and the effectiveness of the vaccine at preventing the flu. To the extent that believing or disbelieving is a time-invariant characteristic, our specifications with person-fixed effects will control for that. We do this with both personal and aggregated (from the claims data) measures of flu incidence and vaccine failure rate. We compare these two measures with the CDC's own measures of flu incidence and effectiveness rate, as well as include year-fixed effects which should control for average annual differences.

Table 2 also reports the incidence of the flu shot, conditional on decisions and outcomes of the previous year. Three broad patterns are evident: people who get a flu shot in one year are more likely to get it the following year; people who get the flu one year are more likely to get a flu shot in the following year; and people who get both are less likely to get a flu shot in the following year than if they did not get flu (but did get a flu shot). This pattern is evident in the raw cross-tabulations, and it will be repeated below in regressions. The magnitudes will vary by specification or by group, but the qualitative pattern is the same. Such robustness suggests that individuals do learn by suffering.

# 4 Who Gets a Shot When?

## 4.1 Main Specification

We now consider the determinants and patterns of individual take-up. Consider the following linear equation, characterized for person $i$, living in region $r$, at time $t$:

$$s_{irt} = \beta_1 s_{irt-1} + \beta_2 f_{irt-1} + \beta_3 s_{irt-1} \times f_{irt-1} + \alpha_i + \tau_t + \epsilon_{irt}, \quad (1)$$

where $s_{irt} = 1$ if person $i$ living in region $r$ got a flu shot in flu season $t$; $f_{irt-1} = 1$ if person $i$ living in region $r$ got a flu shot in flu season $t-1$; $\alpha_i$ is an individual fixed effect, while $\tau_t$ is a flu-season fixed effect. Alternate specifications that do not exploit the panel nature of the data replace the individual fixed effect ($\alpha_i$) with a county- or other geography-specific fixed effect ($\alpha_r$).

A few characteristics of Equation 1 are worth mentioning. First, we are predicting present day outcomes using the previous year's outcomes. Specifications that include individual fixed effects will limit the variation available to identify the $\beta$s. Importantly, this means that they are identified from people who do not always get the flu shot or get the flu. As we will argue, this is likely a differentially selected population.

Second, a great deal of the relevant variation in flu-shot effectiveness will be common across regions, since the effectiveness of the flu shot depends upon the make-up of the flu shot made at the national level. The flu shot inoculates individuals against particular strains of the flu, whether or not those are the prevalent strains of the flu present nationally. This national variation will be picked up in the season fixed effects. Thus, the variation being exploited here is all individual or region-season specific.

The various controls allow for us to exclude certain explanations for the behavior that are not learning. For example, there may be robust differences in public health infrastructure (e.g., education and awareness campaigns, or public health clinics that provide access to the flu shot) across counties, and that may confound some of these patterns. To the extent that those differences do not vary over time, our county- and MSA-fixed effects should account for this. Alternatively, individuals may vary in their risk aversion, and that could influence their

willingness to engage in preventive care. To the extent that risk aversion is fixed for a person over time, the individual fixed effects should account for this. Likewise, static susceptibility to illness or knowledge of health risks is accounted for in the individual fixed effects.

After including these controls, we argue that the correlation between an individual's flu shot take up in year t and her flu shot and flu outcome in $t-1$ is driven by "learning by suffering."' One alternative to the learning explanation is mean reversion in the errors: individuals face idiosyncratic costs to an activity that fluctuate over time. A person might face low costs in one year and take up an activity; in the following year, by simple mean reversion, an individual is likely to receive a higher cost and less likely to engage in that same activity. This person is not "learning" that they dislike something; they are just responding to idiosyncratic shocks. Note that this mean-reversion explanation cannot explain the positive relationship between flu shot in two consecutive years. Thus, in order for mean reversion to explain the had-flu/flu shot interaction coefficient, the mean reversion must systematically impact persons who get the flu conditional on getting a flu shot. We believe this type of conditional error is unlikely.

## 4.2 Baseline Results

The first column of Table 3 reports the baseline estimates of the learning parameters. All specifications include individual fixed effects. There are three coefficients of interest. All are probabilities relative to the omitted category of neither having had the flu shot nor having had the flu in the previous year. This ommited category has a thirty-three percent take-up rate of the flu shot the following year. Take, for instance, the estimate for $\beta_1$, in the "Had shot" row. After accounting for individual-fixed effects, individuals who had the flu shot in the previous year were 3.3 percentage points (or ten percent, against the baseline of 33 percent) more likely to get a flu shot, presuming they did not get the flu the previous year. Had they the flu the previous year ($\beta_2$, "Had flu" row), but not the flu shot the previous year, they would have been 4.8 percentage points more likely to get a flu shot.

The estimates are quite substantial; all correspond to at least a three percentage point change in the underlying likelihood of getting a flu shot, against a baseline population average of thirty percent. The medical advice that nearly all persons should get a flu shot did not

vary over this period; so to the extent that the "rational" decision is to follow medical advice, there should be no variation to identify the point estimates. Even if individuals thought they should "learn" from experience, there is little reason to think that they should adjust their priors so quickly, particularly since these individuals are over 65 years old, and have already had a great deal of exposure to both flu risk and the flu shot.[7]

In Table 3 column one, the sum of the three key coefficients ($\beta_1, \beta_2, \beta_3$) is practically zero, suggesting that persons who get the flu shot and the flu in one season are as likely to have a flu shot the following year as someone who neither had the flu shot nor had the flu in the previous year. When a person is given preventative care and the person still gets sick, it is as though they did not take the preventive care and did not get sick in spite of it.

As mentioned above, respiratory illnesses (any ICD-9 codes between 480 and 490) may be confused for the flu. We replace the flu-related indicators with the broader respiratory-illness category in the second column of Table 3. The same patterns remain, and two of the three grow stronger, particularly the interaction of having had a respiratory illness and having had a flu shot. Thus, this learning-by-suffering corresponds to a broader class of suffering than that specifically diagnosed as the flu. (For the balance of the paper, we return to the more narrow definition of the flu-related illness.)

There are many ways to suffer from the flu, and some may be more severe than others. In order to assess the role of the intensive margin in learning by suffering, we include additional indicators if the flu-related illness involved admission to a hospital. The results for this specification are in the third column of Table 3. The interaction of being admitted for the flu in spite of having the flu shot further decreases the likelihood that the person subsequently received a flu shot, relative to receiving outpatient care for the flu shot. That is, the more suffering, the more intense the learning.

To check the role of geographic variation in the specification with individual fixed effects, we add county-level averages to the fourth column of Table 3. Among the county averages, the only statistically significant coefficient is the previous year's flu shot average against subsequent flu shot take up. We are hesitant to interpret the regional-average coefficients

[7]When we limit our sample to exclude those who always or never got the flu shot while in the sample, the results remain practically the same. This exercise reduces our sample size, but The results for this specification are available upon request.

as evidence of region-wide learning. We view these regional averages as controls for factors that may change over time within a region. Further, adding these regional averages do not change how individuals learn from their own experience of flu and flu shot.

Finally, as we have seen, the 2009 flu season was unusual, due to the Bird flu. The 2009 season is included in these results, but the patterns remain even if we drop 2009.

## 4.3 Does the Learning Attenuate Over Time?

It is theoretically ambiguous if the learning attenuates over time. Assuming the underlying truth of individual susceptibility and vaccine effectiveness does not change over time and the degree of noise remains stable, Bayesian learners should give equal weight to each year's experience. However, if the underlying truth evolves over time, one should give more weight to recent than remote experience.

To test this, Table 4 reports the point estimates for two separate specifications both of which include three years worth of lags, instead of just one. We focus on the 2008 season, since it allows us to consider three years of prior experience, while avoiding the 2009 Bird Flu season (see below). Because only one year of observations are used, individual-fixed effects are not feasible. However, these estimates do account for county-level fixed effects. Thus any county-level variation that drives either flu risk or flu shot take-up is absorbed, and the coefficients reflect the remaining variation.

Results suggest that the effect of prior experience does attenuate over time. The signs and magnitudes reported in Table 4 are comparable to the baseline in Table 3. One explanation is that people believe their individual susceptibility to the flu virus evolves over time, and so does the effectiveness of flu shot. This belief implies that recent experience is more informative. Another possibility is that recent experience is just more salient to individuals when they predict the likelihood of flu or the effectiveness of flu shot.

## 4.4 Did the Bird Flu Outbreak of 2009 Break the Learning Pattern?

The patterns just described are consistent with the notion that individuals are learning over time in a way that privileges their own personal experience. One way to test whether or not individuals can learn about their learning is to assess these patterns before and after a well-established shift in knowledge and expectations.

The Bird Flu (H1N1) Outbreak of the 2009 season provides such an opportunity. The 2009 flu season was distinct from the other seasons, in that the flu oubreaks started earlier than in other years. This is evident in Figure 2, which plots the histogram of office-based claims for flu care by date, separately for 2009 season and other season. According to the CDC's official history of the season, the H1N1 virus represented a genetically distinct strain of the flu. Along several dimensions, this flu season was distinct from previous flu season; to the extent that foretold changes in future seasons, one's ability to learn from previous evidence may wane if the external validity of previous evidence is put into question.

The fifth column of Table 3 reports the point estimates for Equation 1, with interactions for the main estimates of interest with two post-2009 indicators: three interactions for the 2010 season, whose flu and flu shot outcomes are drawn from the 2009 season; and another three interactions for 2011 and 2012 combined. Individual fixed effects are included in this specification. The flu shot outcomes for the 2009 flu season are excluded from the estimation sample, since the timing of the flu season and supply issues were unusual that year. These outcomes are used as explanatory variables for the 2010 outcomes. Flu shot behavior that year may have reflected the general uncertainty of that season and the expanded public health efforts that accompanied it. Because we are interested in changes in behavior before and after that season, it is excluded as an outcome from this particular analysis.

The point estimates reported in Table 3 indicate that the three parameters of interest in Equation 1 were attenuated for the 2010 Bird Flu season, relative to prior patterns. This is particularly true for those indicators involving having had the flu, where the magnitude of the 2010 interaction drives down the 2010 pattern, based from the 2009 results, to nearly zero. However, the interaction between the 2011 and 2012 indicator and those two coefficients of

interest shrinks, demonstrating that this attenuation was short-lived. That is, individuals' response to getting sick in 2010 returned to patterns more similar to the pre-Bird flu seasons. However, the response to having had a flu shot in the previous year, which was practically the same in 2010 vis-a-vis pre-2009 patterns, is itself driven down to near-zero after 2010.

# 5 Measuring the Aggregate Implications

We now consider the aggregate implications of these dynamic patterns. That is to say, how do these patterns influence the incidence of the flu or the flu shot rate? Further, given that the incidence of the flu inherently fluctuates due to environmental factors, such as the transmission rate, the dynamics described above may amplify or mitigate consequences of those fluctuations.

Consider the following Markov transition:

$$\begin{aligned} s_{t+1} &= (1-s_t)(1-f_t)\pi_{00} + (s_t)(1-f_t)\pi_{10} + (1-s_t)(f_t)\pi_{01} + (s_t)(f_t)\pi_{11} \\ f_{t+1} &= s_t \cdot e_t^s + (1-s_t) \cdot e_t^{ns} \end{aligned} \tag{2}$$

where $\pi_{..}$ are the probabilities of getting the flu shot, conditional on getting sick and having had the flu shot in the previous period. We introduce two additional variables: $e_t^s$ measures the incidence of the flu in individuals who do get the flu shot, while $e_t^{ns}$ is the incidence of the flu in people who do not get the flu shot, with $e_t^{ns} > e_t^s$ if the flu shot is a prophylaxis against the flu.[8] Our initial measures of $\pi_{..}$ are taken from Table 2. The first equation is clearly Markov, as it relates state variables in period $t$ to an outcome in $t+1$. The second equation is less obviously so, since it relates state variables in $t$ to an outcome that occurs in period $t$ but is indexed as period $t+1$. Due to the timing (people choose to get the flu shot or not prior to the severity of the flu season, or the preventive value of the flu shot, being realized), it is period $t$'s flu realization that informs the decision whether or not to get the flu shot in period $t+1$. Given the six parameters (the four $\pi$s and the two $e$s), we can calculate a steady state levels of the flu shot and the flu. Note that this may not match any particular

[8]The ratio of the two $e$s is not directly related to the flu shot effectiveness measure used above, since they measure similar, though, distinct things.

year's flu shot and flu rates, since those are consequences of idiosyncratic variation in the parameters.

We first consider how variation in the four $\pi$s alters the steady state values of $s = s_{t+1} = s_t$ and $f = f_{t+1} = f_t$. Figure 3 plots the path to convergence in $s$ and $f$ for the baseline parameter values. It takes five periods for the values of $s_t$ and $f_t$ to converge to within one percent of their steady state values. The steady-state values were calculated by repeated iteration of the transition functions in Equation 2 until the difference between iterations was less than $1 \times 10^{-10}$.

As documented in Table 5, removing, or at least adjusting, the way individuals respond to personal outcomes can substantially change the equilibrium levels of behavior. The first row reports the behavioral parameters and equilibrium outcomes using the conditional means of behavior in Table 2. The second row reports the same, but we remove the learning-by-suffering patterns ($\beta_1$ through $\beta_3$) from the panel estimates of Table 3. We call this parametrization "no learning by suffering."

Removing these "learning" patterns decreases the rate of the flu shot and increases the rate of the flu. The main mechanism through the third column, which reflects the year-on-year rates of individuals who get the flu shot but do not get the flu—the largest group whose patterns are adjusted. The flu shot rate falls four percentage points (over ten percent against a baseline of 35 percentage points), while the flu rate grows one one-hundredth of one percent. Figure 4 plots out the flu incidence and flu shot rates as they transition year-to-year as governed by Equation 2. The convergence to these new equilibrium values appears to take about five years.

The final column of Table 5 maps the new equilibrium flu incidence, $\widetilde{r_t}$, into discounted spending on the flu over a ten year period, relative to the baseline behavior and flu incidence ($r_t$) of the first row.

$$PDV = \sum_{t=0}^{10} \delta^t (\widetilde{r_t} - r_t) c_{flu} POP_{MFFS} \tag{3}$$

The average sum of all medical spending on the flu for a diagnosed case ($c_{flu}$) is approximately \$1,000. This averages across those who receive just outpatient care and those who are admitted to the hospital, and focuses exclusively on the care that is directly linked to the

flu diagnosis. The flu rate is multiplied by $POP_{MFFS}$, 55 million, which is (approximately) the number of Medicare FFS beneficiaries in this time period. We use a one-year discount factor, $\delta$, of 0.98.

The ten-year medical spending on the flu increases by nearly fifty-million dollars were the learning-by-suffering patterns of the flu shot take-up removed. With the fixed effects estimates removed from the transition probabilities (i.e, $\beta$s removed from their corresponding $\pi$s), the flu rate increases because the flu shot rate decreases. Individuals are more likely to get a flu shot in a subsequent season ($\beta_1 > 0$) when the flu shot appears to work—i.e., when the person gets a flu shot and they do not get the flu. A secondary effect ($\beta_2 > 0$) is that people are more likely to get the flu shot when they didn't get a flu shot and got the flu. The magnitude of these effects on the overall spending pattern is proportional to the likelihood of their occurance; thus, the coefficients tied to getting the flu ($\beta_2$ and $\beta_3$) are weighed less heavily in the overall spending patterns since they change the behavior of fewer people.

The third row of Table 5 halves the differences between the raw patterns and those with the learning parameters. That is, the new $\pi$s are $\pi_{\cdot\cdot} - \frac{\beta_{\cdot}}{2}$. The rates of the flu and the flu shout split the difference from the raw patterns and those adjusted by removing the $\beta$s.

We also consider another hypothetical: What do these dynamic patterns tell us about the efficacy of interventions that increase the use of the flu shot. We consider a one-time increase in the flu-shot rate to twenty percent over the steady-state rate for the baseline parameters. That is, an increase from roughly 36 percent to 42 percent. The one-time increase in the flu shot rate, and the associated decrease in the flu rate, deteriorate quickly. Under the full learning (i.e., the raw dynamic patterns observed in the data), the flu shot rate falls roughly two percentage points per year. Correspondingly, the flu rate quickly returns to the steady state value within five periods. Both are plotted at the blue lines in Figure 5. Flat orange lines plot the flu shot rate and flu rate if the one-time increase in the flu shot persisted.

We also apply the "no learning" dynamics, i.e., those that remove the fixed-effects estimates from the raw dynamics, as in the second row from Table 5 . The increase to 42 percent take-up rate is larger than for the full learning specification. The pace at which the flu shot rate falls from this one-time increase is just as stark; the subsequent decreases in the flu shot rate are four to five percent for the next three periods, with a return to the steady

state value after five periods.

The strong attrition from a one-time increase in the flu shot rate can lead to substantial overstatements in the savings associated with (purportedly permanent) increases in the flu shot rate. The (discounted) ten year savings in flu spending from a permanent twenty-percent increase in the full-learning flu shot rate are over $116 million. Those savings fall to $28.5 million after accounting for the full-learning transitions. The no-learning equilibrium flu shot rate is lower, so the cost savings to a permanent increase in the flu shot rate is $179 million. The savings that account for the no-learning transitions is only $39.5 million.

# 6 Conclusion

An annual flu vaccine (or shot) is one of the least controversial and most widely-recommended preventative health measures. The Center for Disease Control's Advisory Committee on Immunization Practices (ACIP) suggests that everyone six months or older get an annual flu vaccine, absent pre-existing allergies or severe health conditions. This advice is consistent with the vaccine's well-established effectiveness at limiting both private and social risks of getting the flu.

We document that the individual decision to take up the flu shot responds strongly to personal, idiosyncratic circumstances. While the medical community urges nearly all individuals to get the flu shot in spite of personal particulars, individual decisions appear to be strongly linked to them. Future advocacy campaigns may well take these patterns into account.

# References


Bodine-Baron, E., R. Varadavas, S. Nowak, and N. Sood (2013): "Conforming and Non-Conforming Peer Effects in Vaccination Decisions," *NBER Working Paper*.

Brenzel, L. E., T. Barham, and J. A. Maluccio (2007): "Beyond 80%: are there new ways of increasing vaccination coverage? Evaluation of CCT programs in Mexico and Nicaragua," in *Evaluation of CCT Programs in Mexico and Nicaragua. iHEA 2007 6th World Congress: Explorations in Health Economics Paper*.

Bronchetti, E. T., D. B. Huffman, and E. Magenheim (2015): "Attention, intentions, and follow-through in preventive health behavior: Field experimental evidence on flu vaccination," *Journal of Economic Behavior & Organization*, 116, 270–291.

Calandrillo, S. P. (2003): "Vanishing vaccinations: why are so many Americans opting out of vaccinating their children," *U. Mich. JL Reform*, 37, 353.

Carman, K. G. and I. Mosca (2011): "Who Takes Advantage of Free Flu Shots? Examining the Effects of an Expansion in Coverage," .

Chang, L. V. (2015): "The Effect of State Insurance Mandates on Infant Immunization Rates," *Health economics*.

Chapman, G. B., M. Li, J. Vietri, Y. Ibuka, D. Thomas, H. Yoon, and A. P. Galvani (2012): "Using game theory to examine incentives in influenza vaccination behavior," *Psychological science*, 0956797612437606.

Dai, T., S.-H. Cho, and F. Zhang (2015): "Contracting for on-time delivery in the US influenza vaccine supply chain," *Manufacturing & Service Operations Management, Forthcoming*.

Engle-Warnick, J., S. Laszlo, N. Mishagina, and E. Strumpf (2013): "Coordination, Common Knowledge and an H1N1 Outbreak," *CIRANO-Scientific Publications 2013s-10*.

Finkelstein, A. (2004): "Static and dynamic effects of health policy: Evidence from the vaccine industry," *Quarterly Journal of Economics*, 119, 527–564.

Fiscella, K., K. Holt, S. Meldrum, and P. Franks (2006): "Disparities in preventive procedures: comparisons of self-report and Medicare claims data," *BMC Health Services Research*, 6, 1.

Gostin, L. O. (2015): "Law, ethics, and public health in the vaccination debates: politics of the measles outbreak," *JAMA*, 313, 1099–1100.

Kahan, D. M., D. Braman, G. L. Cohen, J. Gastil, and P. Slovic (2010): "Who fears the HPV vaccine, who doesnt, and why? An experimental study of the mechanisms of cultural cognition." *Law and human behavior*, 34, 501.

Manski, C. F. (2010): "Vaccination with partial knowledge of external effectiveness," *Proceedings of the National Academy of Sciences*, 107, 3953–3960.

——— (2014): "Vaccine Approvals and Mandates Under Uncertainty: Some Simple Analytics," *NBER Working Paper*.

McConeghy, K. W. (2014): "The Impact of Pharmacy-Based Immunization on Vaccination Rates and Prevalence of Influenza Virus," Ph.D. thesis, University of Illinois at Chicago.

Milkman, K. L., J. Beshears, J. J. Choi, D. Laibson, and B. C. Madrian (2011): "Using implementation intentions prompts to enhance influenza vaccination rates," *Proceedings of the National Academy of Sciences*, 108, 10415–10420.

Moyer, M. W. (2018): "Anti-Vaccine Activists Have Taken Vaccine Science Hostage," *New York Times*.

Naprawa, A. Z. and D. R. Reiss (2014): "Medical Advice and Vaccinating: What Liability?" *Available at SSRN 2495971*.

Opel, D. J., R. Mangione-Smith, J. D. Robinson, J. Heritage, V. DeVere, H. S. Salas, C. Zhou, and J. A. Taylor (2015): "The influence of provider communication

behaviors on parental vaccine acceptance and visit experience," *American journal of public health*, 105, 1998–2004.

Rao, N., M. M. Mobius, and T. Rosenblat (2007): "Social networks and vaccination decisions," .

Reiss, D. R. (2013): "Compensating the victims of failure to vaccinate: what are the options," *Cornell Journal of Law and Public Policy*, 23, 595.

Ryan, B. J. and P. Nourmohammadi (2010): "Can My Boss Stick it to Me? Courts May Give a Shot in the Arm to Employer-Mandated H1N1 Vaccines," *SSRN Working Paper Series*.

Scherer, F. et al. (2007): "An industrial organization perspective on the influenza vaccine shortage," *Managerial and Decision Economics*, 28, 393–405.

Stone, E. G., S. C. Morton, M. E. Hulscher, M. A. Maglione, E. A. Roth, J. M. Grimshaw, B. S. Mittman, L. V. Rubenstein, L. Z. Rubenstein, and P. G. Shekelle (2002): "Interventions that increase use of adult immunization and cancer screening services: a meta-analysis," *Annals of Internal Medicine*, 136, 641–651.

Tsutsui, Y., U. Benzion, S. Shahrabani, and G. Yom Din (2008): "Analyzing the decision to get flu shot: an empirical study," *Available at SSRN 1118243*.

Ward, C. J. (2014): "Influenza vaccination campaigns: is an ounce of prevention worth a pound of cure?" *American Economic Journal: Applied Economics*, 6, 38–72.

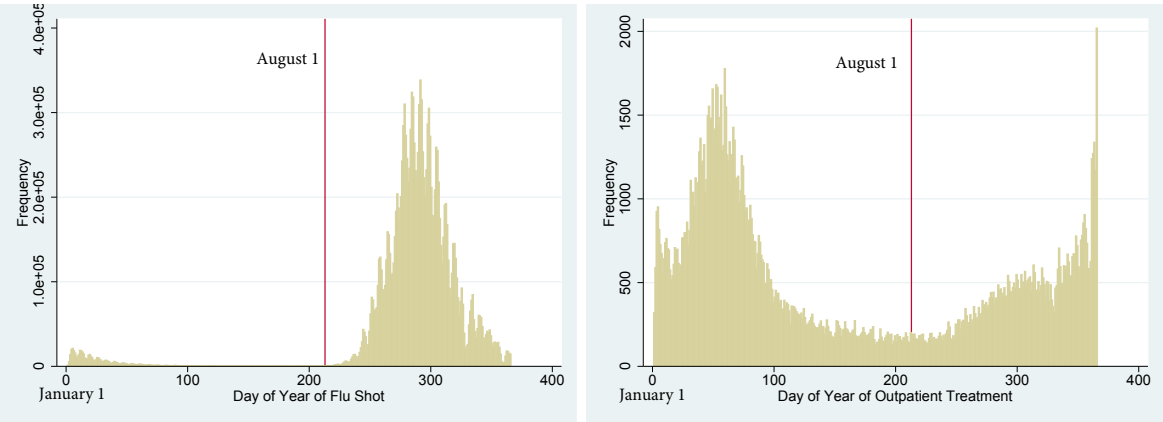


(a) No. flu shots in Medicare FFS 65+ by day of year

(b) No. outpatient flu claims in Medicare FFS 65+ by day of year, offset

Figure 1: The number of flu shot claims and outpatient claims with flu ICD-9 code. The red line on the flu shot graph marks August 1, the start of the flu season. The flu treatment figure shifts the day of year, so that day one is August 1.

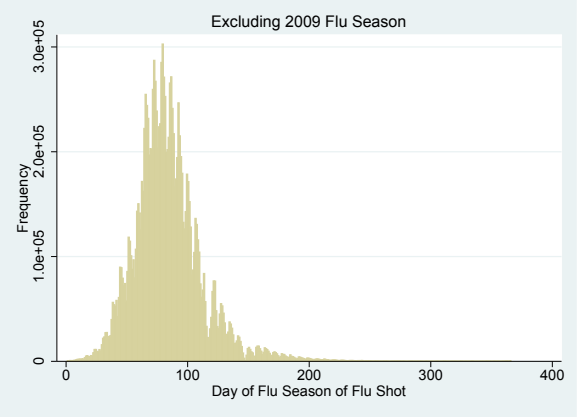


(a) No. flu shots in Medicare FFS 65+ by day of year, 2009

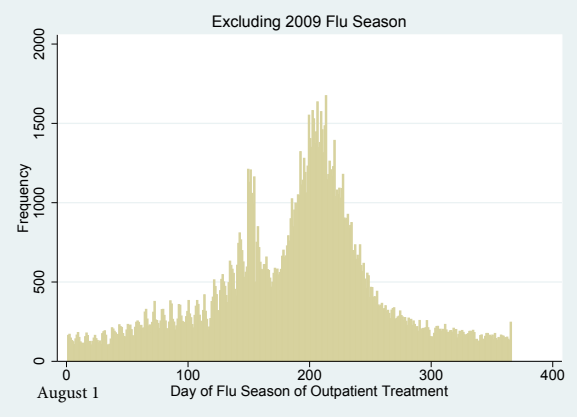


(b) No. outpatient flu claims in Medicare FFS 65+ by day of year, offset, 2009

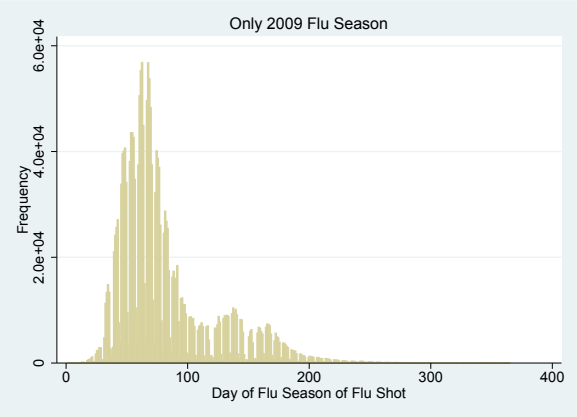


(c) No. outpatient flu claims in Medicare FFS 65+ by day of year, offset, 2009

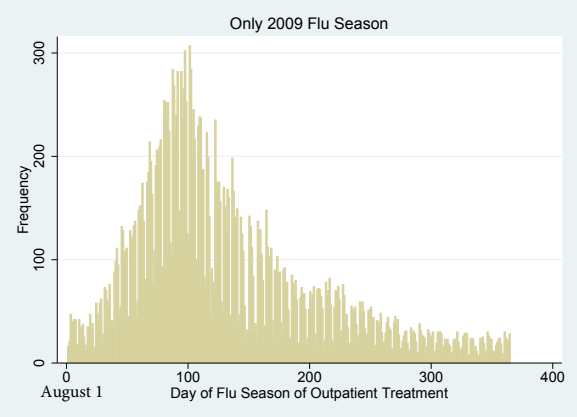


(d) No. outpatient flu claims in Medicare FFS 65+ by day of year, offset, 2009

Figure 2: The number of flu shot claims and outpatient claims with flu ICD-9 code, separtely by 2009 season or not 2009 season. The 2009 flu season was the bird flu season, which began earlier than other seasons. All figures shift the day of year, so that day one is August 1.

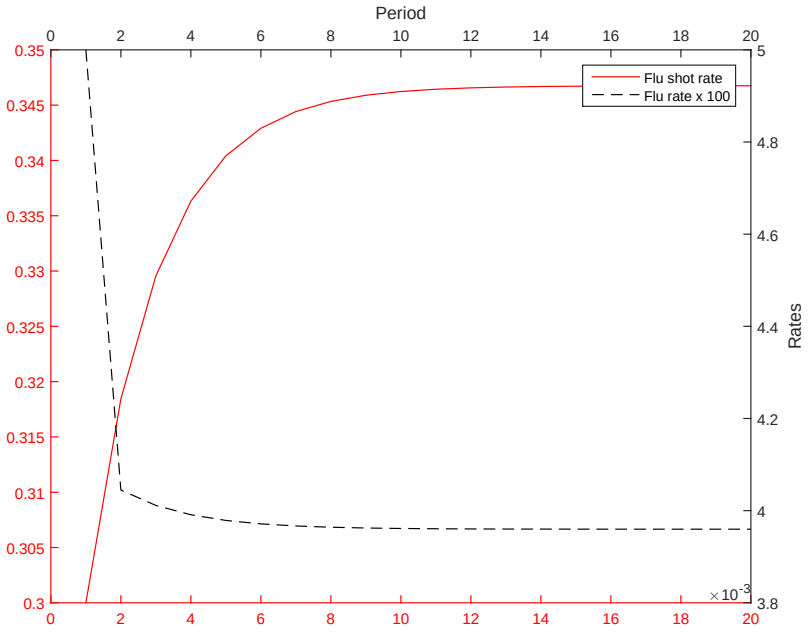


Figure 3: Convergence to steady state flu rate and flu shot rates from initial guess. Calculations made at baseline parameter values.

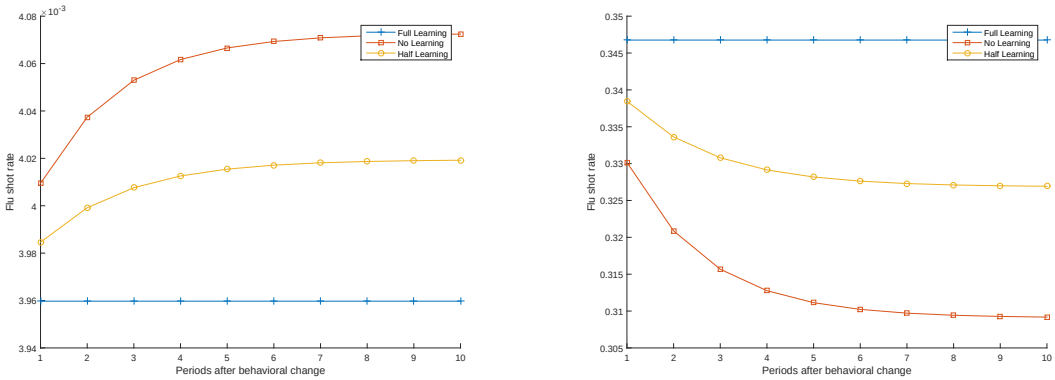


(a) (b)

Figure 4: Convergence to steady state flu rate and flu shot rates for changes to behavioral parameters. Calculations start at full learning equilibrium and the dynamics are then governed by the various changes to learning.

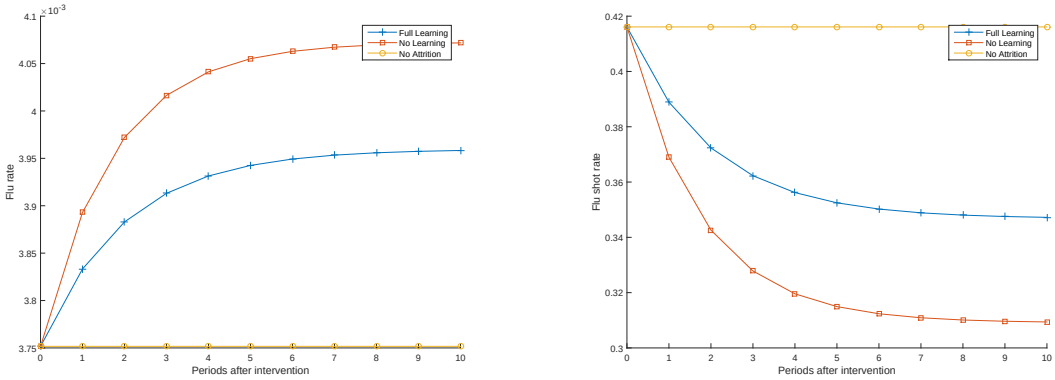


(a) (b)

Figure 5: Divergence from one-time boost in flu rate and flu shot rates, by varying behavioral parameters. Calculations start at one-time twenty-percent increase in the flu shot rate and the dynamics are then governed by the various changes to learning.

Table 1:

| Disease | ICD-9 Code |
|---|---|
| hline Flu | 487.0 |
| | 487.1 |
| | 487.8 |
| Avian Flu | 488.01 |
| | 488.02 |
| | 488.09 |
| H1N1 | 488.11 |
| | 488.12 |
| | 488.19 |
| Novel A | 488.81 |
| | 488.82 |
| | 488.89 |

Table 2:

| | Mean | Flu Vaccination Rate | Flu Illness Rate |
|---|---|---|---|
| All | 100 | 35.29 | 0.43 |
| 1=Men | 40.3 | 32.34 | 0.38 |
| 1=Women | 60.7 | 37.21 | 0.46 |
| 1=White | 80.7 | 38.59 | 0.40 |
| 1=Black | 8.14 | 21.41 | 0.39 |
| | | | |
| 65 <Age< 75 = | 53.3 | 31.27 | 0.37 |
| Age75+ | 46.7 | 39.88 | 0.49 |

| | Flu Shot Rate$_t$ | |
|---|---|---|
| | Had Shot$_{t-1}$ | Did Not$_{t-1}$ |
| Had Flu$_{t-1}$ | 70.04 | 27.22 |
| Did Not$_{t-1}$ | 74.37 | 13.56 |

Notes: Sample means for the 5% Sample of Medicare beneficiaries. Race as characterized by RTI race code. Age as determined at the end of the calendar year. Vaccination rate and flu illness rate are determined by ICD-9 and CPT codes represented on beneficiaries' Medicare FFS claims.

Table 3:

| | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| Previous year outcomes | | | | | |
| Shot x flu | -0.0747*** | -0.145*** | -0.0704*** | -0.0736*** | -0.0869*** |
| | (0.00346) | (0.00114) | (0.00363) | (0.00346) | (0.00512) |
| Had flu | 0.0481*** | 0.0345*** | 0.0470*** | 0.0473*** | 0.0626*** |
| | (0.00244) | (0.000796) | (0.00259) | (0.00244) | (0.00355) |
| Had shot | 0.0334*** | 0.0435*** | 0.0334*** | 0.0296*** | 0.0906*** |
| | (0.000421) | (0.000428) | (0.000421) | (0.000420) | (0.000555) |
| Admitted x shot | | | -0.0407*** | | |
| | | | (0.0100) | | |
| Admitted | | | 0.00774 | | |
| | | | (0.00615) | | |
| Geographic area average | | | | | |
| Shot x flu | | | | -0.00973 | |
| | | | | (0.0767) | |
| Had flu | | | | -0.0128 | |
| | | | | (0.0489) | |
| Had shot | | | | 0.349*** | |
| | | | | (0.00306) | |
| 2010 x | | | | | 0.0543*** |
| Shot x flu | | | | | (0.00952) |
| 2010 x | | | | | -0.0513*** |
| Had flu | | | | | (0.00705) |
| 2010 x | | | | | -0.00933*** |
| Had shot | | | | | (0.000625) |
| 2011-2 x | | | | | -0.00741 |
| Shot x flu | | | | | (0.00888) |
| 2011-2 x | | | | | -0.0149** |
| Had flu | | | | | (0.00614) |
| 2011-2 x | | | | | -0.0903*** |
| Had shot | | | | | (0.000519) |
| Broader flu definition | | x | | | |
| Constant | 0.330*** | 0.329*** | 0.330*** | 0.218*** | 0.324*** |
| | (0.000277) | (0.000280) | (0.000277) | (0.00103) | (0.000290) |
| $N_{i,t}$ | 14,583,592 | 14,583,592 | 14,583,592 | 14,583,592 | 12,507,052 |
| $N_i$ | 2,823,971 | 2,823,971 | 2,823,971 | 2,823,971 | 2,822,076 |

Notes: Estimates for OLS panel models, with robust standard errors clustered at the individual level. All specifications include individual fixed effects. The outcome of interest is whether or not an individual received a flu shot in season $t$, depending upon various individual-level outcomes in season $t-1$. Specification 3 also includes county-level averages of the outcomes for the previous year. $N_{i,t}$ is the number of person-season observations in each specification and $N_i$ is the number of persons observed in the panel. * significant at 10%; ** significant at 5%; *** significant at 1%

Table 4:
Lag Specification, 2008 Season Observations Only

| Year | Flu Shot | Had Flu | Shot x flu |
|---|---|---|---|
| 2007 | 0.0919*** | 0.399*** | -0.116*** |
| | (0.00557) | (0.000973) | (0.00780) |
| 2006 | 0.0477*** | 0.210*** | -0.0608*** |
| | (0.00759) | (0.00101) | (0.0107) |
| 2005 | 0.0363*** | 0.124*** | -0.0425*** |
| | (0.00616) | (0.000898) | (0.00906) |
| Constant | | 0.157 | |
| | | (.) | |
| $N =$ | | 1,775,626 | |
| $R^2$ | | 0.439 | |

Notes: Estimates for OLS models, with robust standard errors clustered at the county level. This table reports the point estimates for one regression, using the 2008 flu shot decision as the outcome of interest. The control variables are whether the person had the flu, had a flu shot, and those two interacted, for each of the previous three years, in addition to county fixed effects. Each column reports the point estimates for a sets coefficient; each row corresponds to a corresponding lag (i.e., the year of that outcome).
* significant at 10%; ** significant at 5%; *** significant at 1%.

Table 5:

| Had Shot? | No | Yes | No | Yes | Steady State Values | | $\Delta$ 10-Year Spend |
|---|---|---|---|---|---|---|---|
| Had Flu? | No | No | Yes | Yes | Had Shot | Had Flu | |
| Baseline | 0.1356 | 0.2722 | 0.7437 | 0.7004 | 35 | 0.396 | |
| No L-b-S | 0.1356 | 0.1022 | 0.6956 | 0.7751 | 31 | 0.408 | \$49,189,701 |
| Half L-B-S | 0.1356 | 0.1872 | 0.71965 | 0.73775 | 33 | 0.402 | \$25,601,375 |

Notes: The relationship between steady state flu shot rate and flu incidence and behavioral parameters and flu risk. All numbers are percentage points. The first column corresponds to the probability of getting a flu shot if in the excluded category (did not get the flu or the flu shot in the previous year). Each row corresponds to changes to the behavioral parameters, by removing (Row 2) or mitigating by half (Row 3) the learning-by suffering (L-b-S).