

# Adverse Health Events and Vaccine Hesitancy

Svenja Miltner\*

Jonatan Riberth<sup>†</sup>

[Job Market Paper]

[CLICK HERE FOR MOST RECENT VERSION](#)

## Abstract

Vaccination forms a cornerstone of public health, combining private protection with reduced transmission to others. Motivated by this, we study how prior adverse health events affect vaccine hesitancy. We use new, nationwide Swedish administrative data that link individual vaccination records, reports of adverse drug reactions, and detailed healthcare utilization across the entire population. We first examine how individuals draw on their past experiences when making new vaccination decisions. To do this, we analyze a severe, well-identified case: narcolepsy, a chronic neurological disease plausibly induced by the 2009–2010 swine-flu vaccine. We find large reductions in COVID-19 vaccination more than a decade later, with spillovers to close family members, suggesting that individuals rely extensively on own experiences when making decision about their health. The effects do not attenuate among those with high health literacy but show some evidence of attenuation with extensive prior healthcare contact. Second, we assess the broader social costs of routine adverse events by studying serious events from all licensed pharmaceuticals. The average effects of experiencing such events on COVID-19 vaccination outcomes are small. While there are some effects for vaccine side effects, there are virtually none for adverse events from other drugs. This highlights that individuals draw on only a narrow set of experiences in future decisions and suggests that the overall social costs of adverse events are likely limited. In the final part of the paper, we show that similar patterns extend to well-established childhood vaccines. This suggests that severe adverse events can erode trust in official safety communication, rather than merely reflecting learning from experience. While the overall impact of routine adverse events on vaccination is limited, rare but severe vaccine-related events can meaningfully lower uptake, providing a cautionary tale for future vaccination campaigns.

**Keywords:** Vaccine hesitancy, Adverse events, Selective memory

**JEL Codes:** D01, I12, I14

---

\*FU, Freie Universität Berlin. E-mail: s.miltner@fu-berlin.de.

<sup>†</sup>IIES, Stockholm University. E-mail: Jonatan.Riberth@iies.su.se.

The authors are grateful to Arash Nekoei, Torsten Persson, David Schönholzer, Anne-Marie Landtblom, Anna Tompsett, Sofia Amaral-Garcia, Anna Dreber-Almenberg, Ronnie Schöb, Natalia Danzer, Peter Mohr, Max Steinhardt, Olof Östergren, Adrian Adermon, Märten Palme, Peter Nilsson and seminar participants at IIES, Stockholm University, Freie Universität Berlin, Paris School of Economics, SWECOV Workshop, SUDSWEC Workshop, Université Libre de Bruxelles, The Workshop on the Economics of Pandemic Preparedness, SHEA/HEFUU Workshop, The annual meeting of Verein für Socialpolitik 2025. This research was conducted within SWECOV—the Swedish Register-based Research Program on COVID-19—hosted by Stockholm University and financed by Riksbankens Jubileumsfond (RIK21-0004).

# 1 Introduction

Recent estimates suggest that vaccines have averted up to 150 million deaths over the past half-century (Shattock et al., 2024), marking a cornerstone in public health advancements. Yet several high-income countries have experienced an alarming resurgence in recurring outbreaks of formerly controlled diseases in the past decades. The inability to contain vaccine-preventable diseases is mainly driven by insufficient demand for cheap and effective vaccination. Notably, in 2019 the World Health Organization declared vaccine hesitancy—“the reluctance or refusal to vaccinate despite the availability of vaccines”—one of the ten major threats to global health (WHO, 2019).

A natural but understudied source of vaccine hesitancy is the perceived risk of adverse events. This channel is important because adverse drug reactions are common. Between 5–10% of all hospitalizations and a significant share of primary care admissions are related to such reactions (Komagamine, 2024; Insani et al., 2021; Pirmohamed et al., 2004). Beyond their immediate costs in the form of medical treatment and illness, adverse events may result in reduced future uptake of vaccines and other medical treatments. Assessing the scope of these indirect social costs is essential for policies aiming to increase vaccination uptake and maintain population immunity.

Using rich administrative records, this paper provides evidence on the effects of experiencing adverse events on future immunization outcomes. While an existing literature links vaccine hesitancy to information gaps, trust in institutions, and peer influence, evidence on how personally experienced adverse drug events shape subsequent demand for health care later in life remains limited. We quantify these long-run behavioral responses by linking individual-level adverse events to subsequent vaccination decisions.

The first part of this paper unpacks how individuals rely on their own experiences to inform subsequent immunization decisions. The decision to vaccinate requires individuals to trade off the expected health benefits, such as milder disease progression, against the risk of adverse events. Perceptions of risks and benefits are, in turn, shaped by individuals’ past experiences. The impact of such experiences on future vaccination decisions depends on how individuals *generalize* across domains; an adverse event from one drug may influence uptake of that drug, related drugs, or medical treatments more broadly. Drawing on insights from cognitive psychology, when facing a decision to vaccinate, individuals use experiences that are selectively retrieved from memory. The primary driver of what experiences are retrieved and shape decisions is the perceived similarity between experiences and outcomes of decisions at hand.

To study how healthcare experiences are retrieved and used across domains, we leverage cases of

narcolepsy—a chronic neurological condition associated with excessive daytime sleepiness and sudden muscle weakness—linked to *Pandemrix*, a vaccine deployed during the 2009–2010 swine-flu pandemic, when roughly 60 percent of the Swedish population was immunized. Epidemiologists later observed an increased incidence of narcolepsy. Subsequent studies corroborated a causal link; an estimated 150–200 individuals developed narcolepsy from Pandemrix. When COVID-19 emerged, it shared key features with the swine flu—both are fast-spreading respiratory diseases for which authorities rapidly introduced vaccines. These similarities provide a setting where previous adverse events are particularly likely to inform healthcare decisions. Drawing on these parallels, we analyze how experiencing this severe adverse event affected vaccination uptake during the COVID-19 pandemic.

Our empirical strategy is based on comparing immunization outcomes among individuals who developed and reported narcolepsy as an adverse event of the vaccine to otherwise similar individuals who received the same vaccine but did not develop narcolepsy. A key element of our analysis is that we also draw comparisons to individuals who were diagnosed with narcolepsy *prior* to the swine flu vaccine campaign and therefore did not have narcolepsy induced by the swine flu vaccine. This difference-in-differences design isolates the effect of experiencing a severe, vaccine-induced adverse event on later vaccine hesitancy.<sup>1</sup>

We find that experiencing a severe adverse event scars individuals with respect to future vaccination uptake—reducing their likelihood of getting vaccinated, the number of doses they take, and delaying the timing of the first dose. In particular, developing narcolepsy after swine-flu vaccination is associated with a 40 percentage point lower likelihood of vaccinating during the COVID-19 pandemic more than ten years later compared to similar individuals. These effects spill over into personal networks, with close family members of vaccine-induced narcolepsy patients being 10 percentage points less likely to vaccinate.

These findings suggest that individuals are not perfectly informed but instead rely heavily on their own experiences to form an assessment about the common risk of experiencing adverse events from the COVID-19 vaccine. By eliciting perceived risks of infection—as measured by self-testing and seeking medical advice during the pandemic—and thus the perceived benefits of preventive vaccination, we show that changes in perceived benefits are unlikely to be driving the results. Partners of individuals with narcolepsy show similar patterns to other family members, suggesting that the results are not driven by beliefs about genetic predisposition.

---

<sup>1</sup>We define vaccine hesitancy as delaying or abstaining from vaccination. Some scholars, however, use the term to describe individuals with persistent ambivalence toward vaccines—regardless of whether they eventually vaccinate (see [Murphy et al. \(2021\)](#); [Larson et al. \(2014, 2022\)](#)). In their terminology, our group would instead be classified as non-compliers.

We find modest evidence of heterogeneity in the effect of exposure to narcolepsy on COVID-19 vaccination by prior healthcare experiences, suggesting that while other experiences matter, the particularly similar and salient event of narcolepsy crowds out other, likely relevant, experiences. In addition, we do not find smaller effects for individuals with high health literacy. While individuals with high literacy should rely more on official information (COVID-19 vaccines being safe), they are equally affected by exposure to narcolepsy. One plausible explanation is that salient experiences crowd out reliable information, even when that information is accessible.

Finally, individuals who have benefited from vaccinations through exposure to communicable but vaccine-preventable diseases, such as influenza, show positive effects on vaccination outcomes, indicating that positive experiences with vaccines also influence uptake.

In the second part of this paper, we assess the overall impact of routinely experienced adverse events on subsequent vaccine uptake by analyzing reported adverse events from all licensed pharmaceuticals as collected through spontaneous reporting systems used for surveillance purposes.<sup>2</sup> These reports are plausibly representative of the full spectrum of routinely occurring adverse events, encompassing adverse events from both novel and well-established drugs.

The credibility of the design rests on two features. First, we restrict attention to events that are serious enough to require medical attention and that are reported by healthcare professionals rather than self-reported by patients. Second, the richness of our data allows us to match individuals not only on demographic and socioeconomic characteristics but also on prior health status, as reflected in medication use and diagnoses. Specifically, we compare vaccination outcomes for individuals who experienced an adverse event with observably similar individuals who received the same drug in the same year but did not experience such an event.

We find large effects for adverse events from vaccines but small effects for other drugs. Experiencing a non-vaccine adverse event in the decade before the COVID-19 pandemic reduces vaccine uptake by just 0.6 percentage points. In contrast, vaccine-related adverse events lower uptake by 7.5 percentage points. These results highlight that individuals draw primarily on a narrow set of similar experiences to inform new decisions. As with narcolepsy, these effects appear long lasting; There is no evidence that only recent adverse events influence the decision to vaccinate against COVID-19. Insights relevant for policy emerge. Although severe adverse events are common, generalization across domains is narrow. As a result, only vaccine-related

---

<sup>2</sup>Most high-income countries operate such reporting systems to detect new safety signals and to gauge the frequency of adverse events after market authorization of a product.

adverse events have a meaningful impact on subsequent vaccination decisions, implying that the overall contribution of adverse events to vaccine hesitancy is limited. However, the results for vaccine-related adverse events is relevant when the private net benefit of vaccination—ignoring potential future uptake effects—is marginal, e.g., for COVID-19 boosters.

In the final part of the paper, we ask whether the patterns we observe for newly introduced vaccines also extend to well established and safe treatments in early childhood. We find that individuals who developed narcolepsy after swine-flu vaccination are less likely to vaccinate their children with established and well-known vaccines. Similar patterns hold for adverse events from other vaccines, but not for adverse events from other types of drugs. Since childhood vaccines are well established, these findings suggest that the reduced uptake cannot be explained by learning from personal experience alone. A more plausible explanation is that trust in official information takes a hit, leading people to rely more on their own experiences when making vaccination decisions. Finally, to assess whether scientific validation matters, we focus on autism diagnoses, a condition once erroneously associated with vaccination against measles despite the absence of any causal link. When older siblings are diagnosed with autism, parents refrain from vaccinating younger children with the MMR-vaccine. This suggests that scientific recognition is not crucial for individuals' perception about the link between treatment and adverse event.

Our paper relates to a literature examining the effects of previous medical scandals and malpractice on future healthcare demand (Martinez-Bravo and Stegmann, 2021; Alsan and Wanamaker, 2017; Lowes and Montero, 2021; Archibong and Annan, 2023). This literature emphasizes long-term deterioration in trust following health scandals, showing large negative effects on demand for healthcare. While we also study negative experiences and healthcare demand broadly, these papers provide examples of quite extreme historical malpractice. Our data and setting allow us to assess the marginal impact on immunization outcomes of an adverse event in a setting that is relevant for future vaccine campaigns and new pharmaceuticals. Furthermore, the administrative data allow us to study how vaccine hesitancy effects spill over in social networks, such as family members and colleagues.

Our paper is also related to a long-standing literature on decision making under risk, and in particular, a growing literature on the lasting impact of previous experiences for behavior (see for instance Malmendier, 2021a). Our primary focus is on a *novel* risk, such as taking the COVID-19 vaccine, where individuals base their decision on a selection of related, previous experiences. A series of articles have studied how salient experiences shape the decision weights individuals assign to different events (Bordalo et al., 2012, 2022, 2024). Building on this literature on selective memory, our study contributes in three ways. First, we

analyze revealed vaccination behavior rather than stated intentions. Second, we link those choices to recorded experiences through comprehensive administrative data on healthcare visits and prescription-drug use rather than self-reported memories or other proxies for experiences. Third, we exploit a severe health shock whose emotional salience plausibly shaped the thinking of those affected, allowing us to quantify its long-term effect on subsequent medical decisions.

Finally, this paper is related to [Oster \(2018\)](#) and [Giulietti et al. \(2023\)](#) who, rather than studying exposure to adverse events from vaccines, focus on how exposure to vaccine-preventable diseases increases the perceived benefits and uptake of medical treatments. In relation to these articles, our primary focus is on negative experiences with healthcare rather than exposure to diseases.

The paper proceeds as follows. Section 2 reviews the swine-flu vaccination campaign and the ensuing narcolepsy episode. Section 3 develops a conceptual framework and testable hypotheses. Section 4 describes the data, and Section 5 explains the empirical strategy. The results unfold through four sections. In Section 7 we document effects for individuals who were affected by narcolepsy, along with heterogeneity and tests of the propositions from the conceptual framework in Section 7.2. In Section 8 we turn to the results for general severe adverse events. In Section 9 and Section 10 we study effects on child vaccine immunization. Section 11 concludes.

## 2 Background

**The Swine flu vaccination campaign** The constant evolution of influenza viruses frequently leads to global transmission chains that occasionally result in severe pandemics. During the 2009–2010 A(H1N1) swine flu pandemic, several vaccines were introduced to the European market early on. The adjuvanted GSK vaccine Pandemrix became the primary, and in some countries including Sweden and Finland, the sole vaccine used to counter the chain of infection. Amid the perceived urgency to quickly achieve high vaccination rates, Sweden launched a mass vaccination campaign unparalleled both in the history of the national health care system and across Europe. Survey estimates suggest that approximately 60% of the adult population was vaccinated, compared to less than 10% in countries like Germany and France that pursued targeted strategies focusing on risk groups, such as the elderly ([Holmberg and Hedberg, 2020](#); [Mereckiene et al., 2012](#)).

To expedite deployment, Pandemrix received market authorization from the European Commission in late September 2009, after only the initial stages of clinical trials. Vaccination in Sweden began shortly thereafter, coordinated by regional health authorities. Vaccination initially targeted priority groups, such as healthcare

workers and individuals at medical risk, before being extended to the general population.

In 2011, health authorities in Finland and Sweden reported a possible link between vaccination with Pandemrix and an elevated risk of developing narcolepsy. Subsequent studies carried out in collaboration with Swedish authorities backed up the initial beliefs (Persson et al., 2014). This link has later been corroborated with evidence from other countries. The current scientific consensus is that Pandemrix indeed caused narcolepsy (Edwards et al., 2019), although developing swine flu itself may also have served as a catalyst by triggering an autoimmune response. Drawing on observed excess incidence, the Swedish Medical Products agency estimates that 150–200 people developed narcolepsy who otherwise would not have developed narcolepsy at that time (Gauffin et al., 2024).

**Legal consequences and public involvement** Patients who developed narcolepsy after vaccination with Pandemrix were initially directed to the Swedish Pharmaceutical Insurance, a program established through an agreement between pharmaceutical companies operating in Sweden, to pursue reimbursement. As part of this program, a maximum of 150 million SEK (15 million USD) was paid out for all injuries attributable to the vaccine, reflecting the fund’s annual cap. In 2016, the government assumed the responsibility of compensating diseased individuals who had not received adequate compensation from the insurance company due to the insurance coverage limit. Individuals received at most 10 million SEK (900k USD) from the insurance company and the government combined.

The general public engaged with the vaccination campaign as it evolved into a controversy over adverse events. In Panel A of Figure B2, we display the overall exposure to the topic illustrated by the number of newspaper articles in Sweden referencing the Swedish word for Narcolepsy (Narkolepsi). Panel B covers public interest as measured by the number of Google searches. Both proxies of public engagement show three major peaks in absolute numbers: (1) In 2011 around the time Pandemrix was announced to be the cause of the increased incidence of narcolepsy, (2) around 2016 when the discussion on reimbursement for narcolepsy patients peaked as the government assumed the responsibility, (3) 2020 when COVID-19 cases were on the rise. Newspapers published the stories of those affected of the narcolepsy episode, documented their endeavors to claim reimbursement, and referred back to the episode when COVID-19 appeared on the public health agenda. These observations are meaningful to our analysis in two ways. First, there was at least minimal exposure and therefore knowledge about the narcolepsy episode in the entire population. Naturally, this general knowledge became more salient in case of personal diagnoses or by the presence of a vaccine-related narcolepsy diagnosis in close networks. Second, it supports the idea that, in fact, the adverse

event of narcolepsy remained in public memory long past the time of episode, which we incorporate in our conceptual framework via the role of past experiences in future health care decisions.

**Narcolepsy** Narcolepsy is a rare, chronic neurological disorder primarily characterized by excessive daytime sleepiness (EDS), which leads to difficulty staying awake and sudden episodes of sleep during everyday activities. There are two main subtypes: Type I narcolepsy, which includes episodes of cataplexy—sudden muscle weakness triggered by strong emotions—and Type II narcolepsy, which does not. Type I narcolepsy accounts for roughly one in three cases [Ortiz et al. \(2025\)](#). Although early reports suggested that Pandemrix-induced cases were overrepresented in the severe form of narcolepsy (type I), later evidence indicates that milder cases were simply identified later ([Stowe et al., 2020](#)). Individuals undergo extensive medical evaluations before receiving a diagnosis, and with increasingly accurate tests—such as spinal fluid analysis to measure orexin levels—misdiagnoses is expected to be rare. There is a possibility that the disease is underreported among individuals with mild symptoms, as we further discuss in section 4.

Medication can be used to manage symptoms, including EDS and cataplexy, with varying efficacy. Yet, the disease still takes a serious toll on everyday life of the diseased and requires permanent lifestyle changes. For example, in our data, among prime age individuals, those with narcolepsy have around 80% lower income than comparable individuals of the same age. Important for our analysis is the fact that the symptoms are not only salient to the affected individuals themselves, but also to those in their social environment—rendering it plausible that individuals in the social networks of the diseased individuals knows that they indeed have narcolepsy.

During the COVID-19 pandemic, there were no official recommendations regarding the COVID-19 vaccine directed toward narcolepsy patients. We can, however, not rule out that individual practitioners may have been either advocating or discouraging the uptake of COVID-19 vaccines. It is probable, that individuals who developed narcolepsy after Pandemrix have a specific autoimmune predisposition, but there is no evidence suggesting they would be systematically more prone to experience adverse events from the mRNA vaccines deployed during the COVID-19 pandemic. The mechanism by which Pandemrix induced narcolepsy is suspected to have involved the immune system mistakenly attacking sleep-regulating cells in regions of the brain that produce orexin, possibly triggered by a flu protein or an adjuvant ([Ahmed et al., 2014](#); [Mahlios et al., 2013](#)). This pathogenetic mechanism is fundamentally different from the mechanisms underlying serious adverse reactions to mRNA vaccines, such as temporary heart inflammation, which may involve a short-lived immune reaction or a rare immune misfire against heart tissue. Given this difference,

the risks associated with the swine flu and COVID-19 vaccines are plausibly unrelated, offering no clear physiological reason for narcolepsy patients to avoid the latter.

Narcolepsy is the only severe adverse event that was prominently discussed to be associated with Pandemrix. There is no evidence to support the idea that Pandemrix would be associated with more adverse events than other drugs, in general. In our data, 0.002% (0.005%) report mild (severe) adverse events of Pandemrix. For comparison 1% (0.01%) of individuals receiving the COVID-19 vaccine reported mild (severe) adverse events to the Swedish Medical Product Agency.

### 3 Conceptual Framework for Immunization Decisions

Healthcare settings often involve high-stakes and emotionally charged decisions, where individuals face information from a range of external sources, including authorities, personal networks, and news outlets. When guidance is technical or conflicting, navigating healthcare can be difficult for laypeople.

A key feature of our setting is that individuals face unprecedented decisions. The onset of the COVID-19 pandemic was characterized by great uncertainty about the scope of infection, the effectiveness of vaccines and the prevalence of adverse events. As a result, when deciding whether to take the COVID-19 vaccine, people had little COVID-19-specific evidence on the risk of infection, vaccine effectiveness, and adverse events. They therefore had to draw on past similar experiences to guide their decision.

A large literature in economics examines how past similar experiences shape choices across settings—such as exposure to high inflation, stock-market shocks (Malmendier and Nagel, 2011; Malmendier, 2021a,b), or consumer demand patterns (Bronnenberg et al., 2012). A consistent finding is that such experiences leave persistent imprints on high-stakes decisions.

Our framework builds on this insight. We assume that memories stem from one’s own medical experiences (for instance, a previous flu shot) or from those shared within personal networks (for example, an anecdote about a sibling’s treatment). When retrieved, these memories compete for attention; decisions depend not on the entire stock of past experiences but on the subset that becomes accessible at the moment of choice. We formalize this idea in a simple conceptual framework of selective memory, inspired by Bordalo et al. (2012, 2022, 2024) and related to the seminal case-based decision theory of Gilboa and Schmeidler (1995).

The decision maker vaccinates if the perceived benefits of vaccinating  $\hat{\pi}_B B$  outweigh the costs of vaccinating  $\hat{\pi}_C C$ .  $B$  includes both individual benefits, but may also include prosocial motives, such as protecting others from infection. We interpret  $\hat{\pi}_B$  as the perceived probability of not developing an infectious

disease and  $\hat{\pi}_C$  as the perceived probability of developing adverse events from the vaccine. The perceived risk of adverse events is formed according to

$$\hat{\pi}_C = (1 - \theta)\pi + \theta\hat{\pi}_E \quad (1)$$

Estimates of  $\hat{\pi}$  are based, on the one hand, on public information about the true probability of severe adverse events,  $\pi$ , and, on the other hand, on  $\hat{\pi}_E$ , which reflects assessments based on personal experiences. While we focus on the perceived risk of developing severe adverse events,  $\hat{\pi}_C$ , we think of  $\hat{\pi}_B$  as also being shaped by experiences. We interpret  $\theta$  as a reduced-form parameter that captures the weight an individual places on personal experience relative to public information. Conceptually,  $\theta$  is increasing in the cost or difficulty of acquiring information about the true risk of adverse events, and in factors such as low trust in healthcare authorities or limited health literacy. For example, individuals with higher trust or better access to official statistics would be characterized by a lower  $\theta$ . If  $\theta = 0$ , personal experiences such as exposure to severe adverse events, are not relevant and individuals base their decision solely on common knowledge available to them. This is the case when the true risk of vaccination is well established and the relevant information is easily accessible at low cost, such as for instance in the case of common children vaccines, which we study in section 9.

By contrast, when public signals are scarce or contradictory, acquiring reliable risk information is costly. In the initial phases of COVID-19 vaccine roll-out, little was known about the occurrence of adverse events on a large scale. In the absence of perfect information on the costs of the vaccine, individuals rely extensively on personal experiences,  $\hat{\pi}_E$  in assessing the risks of adverse events.

Following [Bordalo et al. \(2024\)](#), we conceptualize the decision-maker as having access to a database,  $E$ , of  $N$  experiences that shapes  $\hat{\pi}_E$ . We think of  $E$  as consisting of events within the healthcare domain—either personally experienced or observed through individuals in their social networks. Experiences are binary vectors of features. For example, "Hospitalized" may be one feature of an experience. Let the target event be  $T = \text{"Severe adverse event after COVID-19 vaccination."}$ . Let, furthermore,  $S(e) \in [0, 1]$  measure the similarity between an event  $e$  and the event of developing severe adverse events from the COVID-19 vaccine  $T$ . Similarity increases in the number of shared features between  $e$  and  $T$ . For example, developing adverse events from a COVID-19 vaccine shares many common features with developing narcolepsy from Pandemrix, as both vaccines were rolled out fast during pandemics and countered versions of influenza viruses with comparable symptoms. This is in stark contrast to events of different domains, such as developing an idiopathic heart attack or even events outside the domain of diseases, such as experiencing hospitalization from injuries

sustained in a car accident.

To rely on an experience, a decision maker must retrieve it from memory. We model the retrieval probability of experience  $e \in E$  as

$$r(e) = \frac{S(e)}{\sum_{v \in E} S(v)} \quad (2)$$

The denominator captures interference: the more past events that resemble this target, the more they crowd each other out, reducing the chance that any single experience is retrieved.

Conditional on retrieving the memory of an event, the decision maker uses this event to *simulate* the event in question, such as developing adverse event after vaccination with COVID-19,  $T$ . For a past experience  $e$ , let  $\sigma(e) \in [0, 1]$  denote its simulation value—how strongly  $e$  contributes to constructing the imagined outcome. Just like with recall, the ability of simulating an event is assumed to be increasing in the similarity between two events  $u$  and  $v$   $\sigma(u) \geq \sigma(v)$  if  $S(u) \geq S(v)$ . That is, memories of events that are similar to the COVID-19 vaccine, such as Pandemrix, are predominantly used to simulate the expected outcome of COVID-19 vaccination and thereby inform individual vaccination outcomes. In this way,  $\sigma$  provides a micro-foundation for how individuals use previous experiences to assess risks;  $\sigma$  reflects the idea that more similar experiences are easier to simulate and are given greater weight and, consequently, exert more influence on the vaccination decision. Consequently,  $\sigma$  also captures the idea that experiences that seem more relevant—precisely because they are easier to simulate—are preferentially used; e.g., an individual who experiences vaccine-induced adverse events, rather than other adverse events, will perceive vaccines as riskier than other drugs.

The experience-based estimate  $\hat{\pi}_E$  is defined as

$$\hat{\pi}_E = \sum_{e \in E} r(e)\sigma(e) \quad (3)$$

The perceived risk is the average of each past experience’s simulated risk  $\sigma(e)$  weighted by how likely that memory is to be retrieved  $r(e)$ .

The conceptual framework, where  $\hat{\pi}$  depends on the database  $E$ , the cost of acquiring information  $\theta$ , and officially provided information  $\pi$ , leads to the following testable predictions, which we examine using narcolepsy as a test bed in Section 7.

**Prediction 1** *Individuals with lower health literacy rely more on their own experiences.*

If  $\theta$  can be interpreted as the cost or difficulty of acquiring information about  $\pi$ , we expect individuals with low health literacy have higher  $\theta$  and hence to adjust their behavior more to adding a new experience to

their database. We rely on three proxies for health literacy to test whether individuals rely on experiences to different extent: (i) having a doctor in the family; (ii) paternal cognitive ability; and (iii) years of schooling.

**Prediction 2** *Suppose adding  $e$  and  $e'$  both increases  $\hat{\pi}_E$ . Then adding  $e$  increases  $\hat{\pi}_E$  more than adding  $e'$  if and only if  $S(e) > S(e')$*

That is, the degree to which  $\hat{\pi}_E$  increases—for sufficiently similar experiences—is increasing in similarity. Most striking is the example of the swine flu pandemic and the COVID-19 pandemic, but in general, we test this proposition by studying experiences with adverse events that will differ in similarity to the target event but not necessarily in severity.

In addition, the conceptual framework predicts that additional memories cause interference because they enlarge the decision maker’s memory database with experiences that compete for retrieval:

**Prediction 3**  *$\hat{\pi}_E$  is less sensitive to adding  $e$  when  $E$  is larger.*

A larger number of experiences crowd-out the probability of recall for a particular experience when assessing the probability of  $T$ . The data allow us to directly measure the number of other experiences with the healthcare system. In particular, we test Prediction 3 by considering heterogeneity in the effects on immunization outcomes of exposure to adverse events with respect to number of other experiences as measured by number of healthcare visits and number of drugs taken.

Research in cognitive psychology suggests that other internal factors, such as emotions and identity, connected to memorized events also play a crucial role in the forecasting of decision outcomes. First, emotional events are more likely to be retrieved from memory than neutral events (LaBar and Cabeza, 2006). Second, emotions are encoded with the event and, when the memory is recalled, that affective ‘tag’ is reinstated and amplifies the experience’s impact on the imagined adverse event outcome (For a recent example in economics integrating the role of emotions in memory, see (Ashraf et al., 2024)). While we refrain from modeling emotions explicitly, the severity of the condition and the uncertainty surrounding reimbursement from governmental agencies likely sparked a range of negative emotions. Local newspapers wrote personal and in-depth stories about the affected individuals, often young children and their families, and their struggles against mistrust from both the general public and the government. Because personal relationships increase the emotional exposure to severe medical conditions, such as narcolepsy, we hypothesize that immediate networks of the diagnosed patients carry stronger effects than more distanced members of personal networks. In particular, we would expect direct family members, who first-hand experience the restrictive nature of the disease, to be most affected by the emotional burden of a vaccine related narcolepsy diagnosis.

## 4 Data

### 4.1 Administrative Records

This paper is conducted within the Swedish Register-based Research Program on COVID-19 (SWECOV) and makes use of data provided by the program. Permission to use the data is obtained from Sweden's Ethical Review Authority (permit numbers 2021-02225, 2022-013550-02, 2022-06118-02 and 2024-02342-02).

We combine data from several Swedish administrative sources covering the Swedish population. First, we use data on drugs and healthcare visits from the Swedish National Patient Register and the National Prescribed Drug Register, administrated by the National Board of Health and Welfare. These sources encompass information on all specialist care visits, diagnoses, and drug prescriptions, covering the period 2005–2022. To identify individuals with narcolepsy we use detailed diagnosis and drug codes on narcolepsy diagnoses and associated pharmaceutical treatments during this period. In particular, we access the complete diagnosis codes (ICD-10-SE) for narcolepsy, G47.4A–G47.4X and drugs commonly taken by individuals with narcolepsy. This allows us to pin down individuals with narcolepsy, regardless of whether they developed it from Pandemrix or not. The remaining drug and diagnosis codes are truncated.

The Public Health Agency provides us with data on all COVID-19 vaccinations in Sweden up until March 2023, allowing us to study vaccination decisions throughout the pandemic. The data entails information on the date of each administered vaccine dose as well as the brand and manufacturer of the vaccine, which allows us to construct our main outcome variables of COVID-19 vaccine hesitancy. Furthermore, The Public Health Agency provides us with novel data from the Swedish Children's vaccination program between 2013 and 2024. We use data on measles, pneumococcus, and diphtheria vaccinations administered during the first two years of life, thereby covering all vaccinations given in early childhood. We also use phone calls to 1177, a medical advice helpline run by the Swedish regional healthcare authorities, from which we observe phone calls related to COVID-19 between 2019 and 2023. We access information on COVID-19 tests that were administered by the public healthcare system during the initial period of the pandemic. We use the data on tests and phone calls to elicit perceived risks associated with developing COVID-19.

The Swedish Medical Products Agency supplies records of every suspected adverse drug reaction reported through the national spontaneous-reporting system between 2005 and 2024. Each report includes date of onset, reporter category (health-care professional or lay person), the suspected medicinal product, MedDRA-coded reaction terms, and the regulatory seriousness. All suspected physiological reactions to medication that are

secondary to its therapeutic purpose are expected to be reported to the system, however under-reporting is known to be substantial.

To build our focal sample of individuals who received Pandemrix, we use registers maintained by the regional healthcare authorities, which are sole holders of this data. We gather vaccination data from 10 of the 21 healthcare regions (see Table D2 for a breakdown of the different regions along with vaccination rates). Out of the remaining 11 regions, nine regions did no longer have access to the data in an accessible format and two regions were not willing to provide the data. In total, data on Pandemrix vaccinations from regional healthcare authorities cover 13% of Sweden's population in 2010.<sup>3</sup>

We use standard registers administrated by Statistics Sweden on socioeconomic and demographic characteristics, such as occupation, income, family links, and place of residence. Apart from defining covariates, we use this data to define social networks that are used to analyze how vaccine hesitancy spreads beyond the diseased individuals. Finally, we access data from the Swedish Military Archives on scores from draft screening tests that individuals conducted between 1979 and 1997, a period when military enlistment was mandatory for men. We use this data to derive a measure of *Cognitive Ability*.

## 4.2 Main Samples & Key Variables

We define individuals in our treatment group as having Pandemrix-induced narcolepsy if they (i) reported narcolepsy as an adverse event with Pandemrix as the suspected drug (ii) are diagnosed with narcolepsy in specialized care at least once after November 1st, 2009, when vaccination with Pandemrix began, but not before. We consider all individuals who received the diagnosis G47.4 (ICD-10-SE) and hence do not make any restrictions on the type of narcolepsy developed. This yields a sample of 346 individuals (of whom roughly half are estimated excess cases.). Panel A of Figure 1 displays the distribution of age at the time of first narcolepsy diagnosis, individuals were on average 17 years old when they were diagnosed with narcolepsy. Most of them report narcolepsy symptoms within three years past vaccination. Based on excess prevalence calculations, about half of the individuals in this sample would have developed narcolepsy even if they abstained from vaccination with Pandemrix. The pathogenesis of narcolepsy is complex and it is impossible to isolate the individual cause of the disease; It is likely that all patients in the treatment group

---

<sup>3</sup>Previous, similar, efforts to collect individual level Pandemrix data was done when the link between Pandemrix and narcolepsy was established. In particular, Persson et al. (2014) manage to collect data for 3.3M vaccinated individuals from seven healthcare regions. While we manage to collect data from a few healthcare regions not previously considered, we are unable to obtain information from the major regions of Västra Götaland, Stockholm, and Skåne.

attribute their disease in parts to Pandemrix. We observe a slight increase in narcolepsy diagnoses following the swine flu pandemic also among individuals who do not report it as an adverse event. These are likely individuals who had some mild form of narcolepsy before the pandemic but did not associate their particular symptoms with a specific health condition, partly because narcolepsy is so rare that even general practitioners may not have been familiar with its specific symptomatology before it became salient in the media.

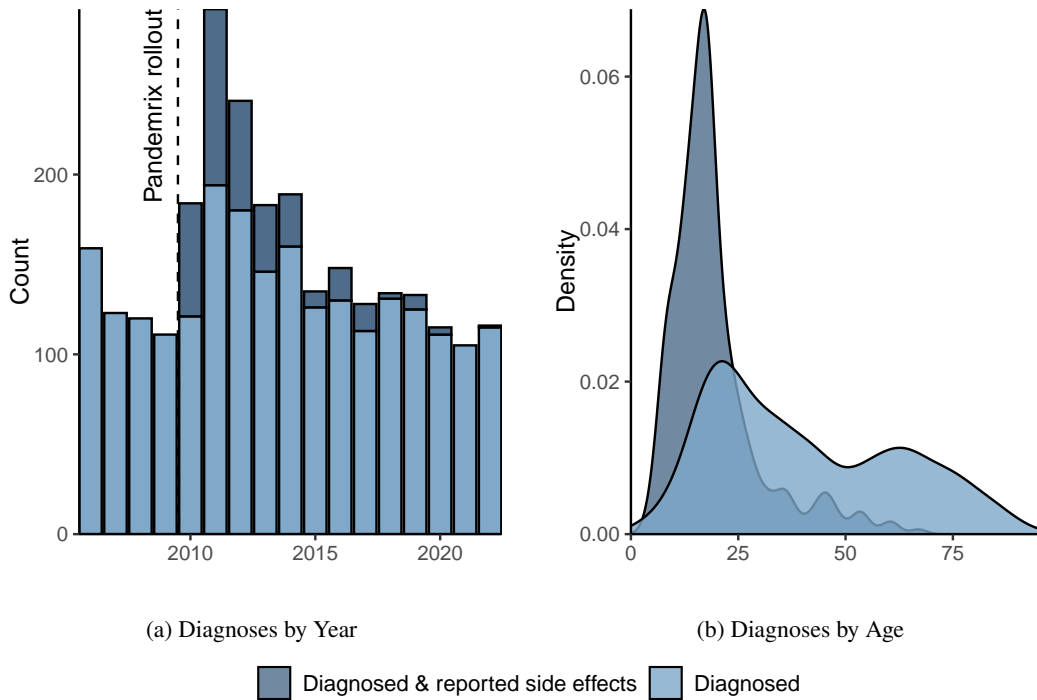


Figure 1: Timing and Age Distribution of Narcolepsy Diagnoses, by Reporting Status

*Notes:* This figure displays year and age of diagnosis for individuals that reported narcolepsy as an adverse event to the Medical Products Agency in dark blue and those that were diagnosed but did not report it as an adverse event in light blue.

**Panel (a)** Count of individuals receiving a first narcolepsy diagnosis in specialized healthcare by calendar year. **Panel (b)** Density of age at first narcolepsy diagnosis.

We define a *pre-swine flu* sample of individuals, who are diagnosed with narcolepsy between January 1st 2005 and October 1st 2009—just before vaccination with Pandemrix begun. We impose the additional restriction that these individuals did not report any other adverse events from Pandemrix. This placebo control group consists of 847 individuals who developed narcolepsy prior to the swine flu pandemic and could, therefore, not have developed it from Pandemrix. Since a diagnosis is commonly recorded when an individual

visits the healthcare system for symptoms related to the disease at hand, this sample also includes individuals who developed narcolepsy much earlier than 2005.<sup>4</sup> These individuals are on average 37 years older than those with Pandemrix-induced narcolepsy.

To assess potential differences in clinical manifestations and severity between narcolepsy patients in the treatment and placebo control groups, we compare their medication uptake. Figure B3 plots prescriptions for common drugs used to treat narcolepsy symptoms in both groups. Consistent with Gauffin et al. (2022), many patients in both groups are prescribed stimulants such as modafinil or methylphenidate for excessive daytime sleepiness, often combined with antidepressants like venlafaxine to manage cataplexies and assist REM sleep. We find no systematic difference in antidepressant use between groups. Prescriptions of dexamphetamine and lisdexamphetamine increase after 2014 in both samples, though somewhat more among the post–swine flu patients, likely because those diagnosed earlier continued established amphetamine routines. Drug schedules thus differ slightly across cohorts, partly reflecting gradual adoption of new treatments. For instance, sodium oxybate, used for cataplexies, became available only in 2012, allowing post–Pandemrix patients immediate access while earlier patients switched later or maintained existing medication. Age differences may also contribute, as sodium oxybate is prescribed with caution due to respiratory risks. Overall, prescription patterns suggest no clear evidence of more severe narcolepsy manifestations in either group.

**Measures of vaccine hesitancy** We consider vaccine hesitancy along three different margins (i) whether an individual took at least one COVID-19 vaccine dose, (ii) the number of doses conditional on taking at least one dose, and (iii) the time elapsed between the dose becoming available to an individual and timing of the first dose. We restrict all analyses to individuals living in Sweden in 2021. Accordingly, we define an individual as unvaccinated if they lived in Sweden in 2021 and do not have a registered COVID-19 vaccination dose. We take on a data-driven approach to define a date when the vaccine is first available to an individual where a date of availability for each combination of birth year and healthcare region is defined. For each birth-year  $\times$  region cell, let  $f(i)$  denote the vaccination-date of the  $i$ -th individual to be vaccinated (Excluding healthcare workers). We define the date of availability as  $f(i^*)$  where  $i^* = \arg \min_{i \leq N-50} \{ f(i+50) - f(i) \}$ . That is, the date of availability is the date that minimizes the number of days elapsed between individual  $i$ 's vaccination and the vaccination date for individual  $i+50$ . In Figure B8 we display the distribution of first vaccinations across time along with defined first date of availability for a sample of 16 region  $\times$  birth year-cells. In general, the distributions of first vaccinations within these region  $\times$  birth year-cells are unimodal and concentrated.

---

<sup>4</sup>It is estimated that around 4 000 individuals in Sweden suffer from some type of narcolepsy (Gauffin et al., 2022).

**Personal networks** We differentiate between close and extended networks. Close networks consist of direct *family*, including biological parents and full siblings of the focal individuals, as well as *extended family*, including biological cousins, uncles, aunts and grandparents of the diseased individual. We consider three extended networks: *neighbors*, defined as individuals living in the same 250×250 m cell as focal sample individuals in 2011; *schoolmates*, defined as individuals attending the same school as the focal individual in 2011; *colleagues*, defined as individuals working at the same plant or establishment in 2021 who are not family members of the treated or control individuals<sup>5</sup>; and *partners*, defined as those who are legally married to or imputed as partners of the focal individual according to Statistics Sweden at any point between 2019 and 2021. If an individual belongs to at least one network of a treated individual, the network member is classified as treated. We assign each network member to only one individual in the focal sample, with ties resolved by random assignment.

**Classification of Adverse Events** Throughout this paper, we classify adverse events into three mutually exclusive categories: Pandemrix-induced narcolepsy, as well as recurring and representative vaccine and non-vaccine adverse events. We restrict attention to adverse events classified as severe—meaning that symptoms were serious enough to require medical evaluation or treatment—and reported by medical practitioners rather than by the affected individuals themselves. We display the most commonly reported severe adverse events in Table B4.

## 5 Empirical Design

Our setting comes with two main empirical challenges. The first challenge—primarily relevant to the analysis of individuals with narcolepsy and discussed in Section 5.1—is to disentangle the effect of experiencing a drug-induced adverse event from the downstream consequences of living with the resulting condition. In other words, we aim to separate the direct effect of exposure to an adverse event from the indirect effect that operates through having narcolepsy itself—for example, through altered interactions with the healthcare system or changes in social networks that shape vaccination attitudes. The second challenge, which is relevant for the analysis of exposure to adverse events in general, relates to the fact that we observe reported adverse

---

<sup>5</sup>We focus on colleagues in 2021, as few affected individuals had colleagues already in 2011. We take the view that narcolepsy was likely salient for colleagues both when the condition first developed and during the COVID-19 pandemic. Colleagues may have observed symptoms in daily interactions, and the Pandemrix– narcolepsy episode was likely still salient when deciding whether to receive the COVID-19 vaccine.

events rather than actual adverse events. In Section 5.2 we discuss under what conditions the differences between reporters and non-reporters may be interpreted as the average effect of developing adverse events. Unsurprisingly, our method will boil down to an assumption about developing and reporting narcolepsy being conditionally independent of vaccine hesitancy. To approximate conditional independence, we match individuals on observables that jointly predict both vaccine hesitancy and narcolepsy onset or reporting. We describe this procedure in Section 5.3.

## 5.1 Isolating the hesitancy-effect

We are interested in the effects on immunization outcomes of developing narcolepsy that may have been caused by Pandemrix. A natural estimator for this effect is

$$\tau_{ideal} = \mathbb{E} [Y_{ith}|t = \text{Pandemrix}, h = \text{Narc. after}] - \mathbb{E} [Y_{ith}|t = \text{Pandemrix}, h = \text{Narc before}]$$

That is, we would ideally compare individuals who develop narcolepsy from Pandemrix to individuals who took Pandemrix and developed narcolepsy prior to Pandemrix vaccination, and who hence could not have developed it from Pandemrix. Because narcolepsy is rare, the overlap between individuals diagnosed with narcolepsy and those observed receiving Pandemrix is small; there are too few cases of individuals who both develop narcolepsy and subsequently receive Pandemrix. Instead, we construct a difference-in-difference estimator that allows us to rely on all narcolepsy cases prior to the swine-flu pandemic. Expanding the previous expression yields

$$\begin{aligned} \tau_{ideal} = & (\mathbb{E} [Y_{ith}|t = \text{Pandemrix}, h = \text{Narc. after}] - \mathbb{E} [Y_{ith}|t = \text{Pandemrix}, h = \text{No narc.}]) \\ & - (\mathbb{E} [Y_{ith}|t = \text{Pandemrix}, h = \text{Narc. before}] - \mathbb{E} [Y_{ith}|t = \text{Pandemrix}, h = \text{No narc.}]) \end{aligned}$$

The first component is the difference between individuals who develop narcolepsy and individuals who take Pandemrix and do not develop narcolepsy. The second component is the difference between individuals who develop narcolepsy *before* the swine flu pandemic and individuals who take Pandemrix and do not develop narcolepsy. Once again, the issue is that the sample of individuals with narcolepsy who received Pandemrix is too small to allow for precise estimation of  $\mathbb{E} [Y_{ith}|t = \text{Pandemrix}, h = \text{Narc. before}]$ . Instead, the key, parallel trends-style, assumption is the following

**Assumption 1** *Constant narcolepsy effect*

$$\begin{aligned} & \mathbb{E} [Y_{iht}|t = \text{Pandemrix}, h = \text{Narc. before}] - \mathbb{E} [Y_{iht}|t = \text{Pandemrix}, h = \text{No narc.}] \\ & = \mathbb{E} [Y_{ith}|h = \text{Narc. before}] - \mathbb{E} [Y_{ith}|h = \text{No narc.}] \end{aligned}$$

That is, the difference in later vaccination outcomes between individuals with pre-existing narcolepsy and individuals without narcolepsy is the same among Pandemrix-takers as in the broader sample.

This assumption allows us to estimate the following object

$$\begin{aligned} \tau_{dd} = & \overbrace{(\mathbb{E} [Y_{iht}|h = \text{Narc. after}, t = \text{Pandemrix}] - \mathbb{E} [Y_{iht}|h = \text{No narc.}, t = \text{Pandemrix}])}^{\tau_{post}} \\ & - \underbrace{(\mathbb{E} [Y_{ith}|h = \text{Narc. before}] - \mathbb{E} [Y_{ith}|h = \text{No narc.}])}_{\tau_{pre}} \end{aligned} \quad (4)$$

Intuitively, we compare vaccinations of individuals who take Pandemrix and develop narcolepsy to individuals who take Pandemrix, filtering out any potential effects of having narcolepsy that is not Pandemrix-induced.

The propensity to take the COVID-19 vaccine among individuals who developed narcolepsy before the swine flu pandemic may be affected in two distinct ways. First, through *biological* consequences of narcolepsy—such as effects on educational and occupational trajectories that shape vaccination behavior. Second, through *informational spillovers*: they may be more hesitant because they are particularly aware of the narcolepsy episode, an awareness that is itself a consequence of the episode.<sup>6</sup> These two mechanisms have different implications for identification.

If  $\tau_{pre}$  primarily captures the biological consequences of narcolepsy, it provides a clean baseline and  $\tau_{dd}$  is a valid estimate of  $\tau_{ideal}$ . However,  $\tau_{pre}$  may also absorb informational spillovers on top of any biological effect, overstating the biological baseline and causing  $\tau_{dd}$  to understate the true causal effect. If spillovers are large relative to the biological effect,  $\tau_{post}$  is the more valid estimate of the causal effect, under the following assumption:

**Assumption 2** *Conditional independence*

$$\mathbb{E} [Y_i(0) | h_i = \text{Narc. after}, t_i = \text{Pandemrix}, X_i] = \mathbb{E} [Y_i(0) | h_i = \text{No Narc.}, t_i = \text{Pandemrix}, X_i] \quad (5)$$

where  $Y(0)$  denotes the immunization outcome that would be observed if individual  $i$  had not developed narcolepsy. That is, conditional on a set  $X_i$  of observable characteristics, developing narcolepsy is as good as random within the Pandemrix group. We return to the plausibility of this assumption in Section 5.3.

---

<sup>6</sup>More generally, this violation of SUTVA also applies, though to a lesser extent, to Pandemrix-takers without narcolepsy (the control group in the post-swine flu sample) as well as to the general population.

It is challenging to disentangle the two mechanisms. We make progress on identifying which is more plausible by considering other severe, chronic, and neurological diseases. Figure 2 displays differences in vaccination rates between individuals with a set of other diseases, controlling for age and gender, along with the estimate for developing narcolepsy prior to the swine flu pandemic,  $\tau_{pre}$ .

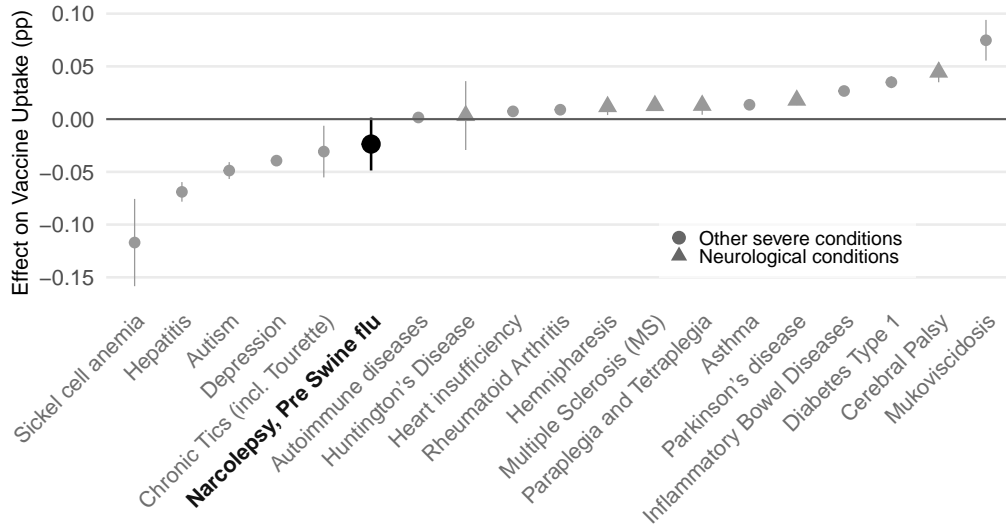


Figure 2: Saliency of the Pandemrix–Narcolepsy Episode

*Notes:* This figure displays estimated differences in COVID-19 vaccination rates between groups of individuals with different diseases. The control group for all regressions is a random subset of 1,000,000 individuals living in Sweden in 2021. All specifications include fixed effects for birth year and gender. *Narcolepsy, pre swine flu* refers to individuals that developed narcolepsy before vaccination with Pandemrix began.

Individuals who developed narcolepsy before the swine flu pandemic have a lower COVID-19 vaccination rate than those with other similarly severe diseases. We cannot rule out that something specific to narcolepsy hinders vaccination—for example, constraints from daytime sleepiness or medical advice. The results, however, support the view that these individuals are primarily influenced by greater awareness of the narcolepsy episode and the perceived risk of severe adverse events, rather than by viewing narcolepsy as a contraindication to vaccination.

## 5.2 Selection into Reporting

We wish to estimate the causal effect of developing an adverse event, such as narcolepsy, on vaccination. We, however, estimate a parameter combining the development *and* reporting of adverse events. Since reporting is

not subject to manipulation, this parameter does not correspond to a well-defined causal effect.

Consider an individual who may develop a drug adverse event. Let  $W_i = 1$  if an individual truly develops an adverse event and 0 otherwise. In general  $W_i$  is unobserved to both the individual and the econometrician. Instead, we observe reporting status:  $D_i = 1$  if  $i$  reports an adverse event. Let  $Y_i(w)$  denote potential vaccination outcomes. We abstract from any additional effect of reporting itself on later vaccination behavior. To the extent that reporting validates or amplifies the perceived link between treatment and symptom, the observed contrast by reporting status may combine the effect of the adverse event with the effect of formal recognition.

We estimate the average difference among reporters and non-reporters:

$$\Delta_D(X) = \mathbb{E}[Y \mid D = 1, X] - \mathbb{E}[Y \mid D = 0, X].$$

While as the object of interest is the average difference between those that develop an adverse events and those that do not:

$$\Delta_W(X) = \mathbb{E}[Y(1) \mid W = 1, X] - \mathbb{E}[Y(0) \mid W = 0, X].$$

The goal is to connect  $\Delta_W(X)$  to  $\Delta_D(X)$ . We invoke the following assumption:

**Assumption 3** *Source irrelevance*

$$\mathbb{E}[Y(0) \mid W = 0, D = 1, X] = \mathbb{E}[Y(1) \mid W = 1, D = 1, X].$$

Among reporters, subsequent vaccination outcomes do not depend on whether the condition was truly an adverse events. The idea is that once a condition is perceived and recorded as a possible adverse event, individuals generally cannot tell whether it was genuinely treatment-induced or would have arisen anyway. Furthermore, we make use of the following assumptions

**Assumption 4** *Missing-at-Random among developers and non-developers (MAR)*

$$D_i \perp Y(1) \mid W = 1, X, \quad D_i \perp Y(0) \mid W = 0, X.$$

Conditional on  $X_i$ , among individuals who develop the adverse event, whether the case is formally reported is independent of future vaccine hesitancy. This is plausible in our setting because the events we study are severe enough to require healthcare and are typically reported by healthcare professionals rather than by patients themselves. The analogous assumption for non-developers is that, conditional on  $X_i$ , mistaken reporting is likewise unrelated to future vaccine hesitancy.

**Assumption 5** *Rare adverse events*  $\Pr(W = 1 \mid X) \ll \Pr(D = 0 \mid X)$

The mass of individuals developing adverse events is small relative to the mass of non-reporters.

Source irrelevance together with MAR among directly give us

$$\mathbb{E}[Y | D = 1, X] = \mathbb{E}[Y(1) | W = 1, D = 1, X] = \mathbb{E}[Y(1) | W = 1, X]$$

Rare adverse events together with MAR among non-developers give us

$$\mathbb{E}[Y | D = 0, X] \approx \mathbb{E}[Y | D = 0, W = 0, X] = \mathbb{E}[Y | W = 0, X]$$

such that

$$\Delta_D(X) \approx \Delta_W(X)$$

We provide a brief mathematical derivation in Section G. In summary, we need four ingredients to recover the causal effect of truly developing the adverse event. First, among individuals who truly experience the adverse event, the decision to report is ignorable once we condition on observed characteristics; some have doctors report it, whereas others do not. Second, the same ignorability must hold for those who did not develop the adverse event: conditional on observables, whether their doctor mistakenly attribute their symptoms to the drug is likewise random. Third, conditional on reporting, the subsequent outcome effect is assumed to be the same for true developers and mistaken reporters—the intuition being that once a symptom is perceived as drug-related, the underlying biological status does not differentially influences future vaccine hesitancy. For instance, a person who would have developed narcolepsy regardless of vaccination reacts no differently than someone whose narcolepsy was vaccine-induced, because neither can disentangle the true source of the condition. Finally, the pool of non-reporters consists almost entirely of genuine non-developers, with only a negligible share of false negatives. Taken together, these assumptions ensure that the average difference in outcomes between reporters and non-reporters recovers the causal effect of actually developing an adverse event under the assumption that developing side effects is conditionally independent of vaccination.

### 5.3 Balance and matching

As is evident from the previous subsection, our empirical approach relies on fairly strong independence assumptions about the data-generating process. The primary concern is that affected individuals are selected in terms of (i) developing and being diagnosed with narcolepsy, and (ii) and reporting narcolepsy as an adverse event. We discuss these concerns in turn.

**Characteristics of individuals with narcolepsy diagnoses** The sharp rise in public awareness after the Pandemrix controversy, together with the fact that diseased individuals require medical treatment, makes it improbable that only a selective subset of patients sought care and received a diagnosis. We quantify how individuals with narcolepsy differ from the general population using extreme gradient boosting (XGBoost), a machine-learning algorithm that sequentially adds decision trees. In particular, we fit a model predicting first-time narcolepsy diagnoses after 2016 using socioeconomic characteristics, and pre-diagnosis health history. The time period set five years after the swine flu pandemic reduces the chance that cases stem from Pandemrix and instead isolates factors linked to developing and being diagnosed with non-vaccine-induced narcolepsy while also letting us observe health conditions before onset. Model performance is summarized by the area under the receiver-operating-characteristic curve (AUC), which ranges from 0.5 (no predictive power) to 1 (perfect prediction). The model is estimated on a subsample that is balanced in terms of birth year.<sup>7</sup> The AUC is therefore interpreted as a measure of the ability to predict who is treated, above and beyond birth year.

As shown in Table B1, this exercise yields an out-of-sample AUC of about 0.55, implying limited correlation between developing narcolepsy and any set of observable socioeconomic and health characteristics. We benchmark these numbers against an unrelated health condition, namely the incidence of a heart attack (I20, I21, I50, ICD-10-SE), which is known to be associated with socioeconomic characteristics as well as individual morbidity profiles (Adhikary et al., 2022). Fitting a model including personal socioeconomic and health history on a sample of 365,000 individuals who developed a heart attack after 2015 combined with a random sample of individuals residing in Sweden in 2021 information yields a relatively higher AUC of 0.61. Taken together, health and socioeconomic characteristics have very limited predictive power for developing narcolepsy.

**Non-random reporting of adverse events** The second challenge is selection in reporting narcolepsy as an adverse event. Importantly, patients had financial incentives to reporting narcolepsy as an adverse event to meet the criteria for the reception of government and insurance reimbursement.

Yet, the remaining main concern is latent vaccine hesitancy, orthogonal to observable socioeconomic and health characteristics, that is correlated with being diagnosed with and reporting narcolepsy as well as vaccination behavior during the COVID-19 pandemic.

We proceed by quantifying how different treated individuals are in terms of observable characteristics compared to the control group. Figure 3, Panel (a), displays differences in a selection of covariates between

---

<sup>7</sup>Specifically, each treated unit is matched to 100 untreated units with the same birth year.

the diseased individuals and the control individuals, conditioning only on birth year.<sup>8</sup> Individuals in the treatment group have higher income and years of schooling. There is no clear pattern regarding their health status, on the one hand they have slightly fewer drug prescriptions and less sick leave but they also make more healthcare visits. Diseased individuals live in more urban areas, which is an artifact of the fact that we do not observe control individuals from the larger metropolitan areas. In Figure B5 we display the same coefficients for the pre-sample of individuals who were diagnosed with narcolepsy prior to the swine flu pandemic compared to individuals of the general Swedish population, independent of Pandemrix status. The differences between these individuals and the general population are similar to the differences in Figure 3. For example, individuals who developed narcolepsy before the swine flu pandemic have parents with more years of schooling than the general population.

**Matching and inference** We use matching techniques to identify individuals that are comparable to those with Pandemrix-induced narcolepsy (treatment group) and to those that developed narcolepsy prior to the swine flu pandemic (pre-swine flu narcolepsy group). In our preferred method we use 1:1 propensity score matching without replacement, exact matching on year of birth and gender, and propensity scores computed using logistic regression. We consider two alternative matching procedures, namely propensity score matching using scores estimated via gradient boosting, and coarsened exact matching. The former aims at better handling interactions, accommodating functional forms that are not well captured by the logit model. Another, practical, advantage of gradient boosting is that it utilizes the full sample and internally accounts for missing values. Consequently, this approach eliminates the need for researcher-imposed decisions regarding the treatment of missing data. Coarsened exact matching aims at dealing with the fact that we have few treated units, making the estimated propensity scores from the the logit model unstable. Instead, matching is done exactly on gender, birth year and coarsened versions of the continuous covariates: Parents' years of schooling, parents' income, parents' number of healthcare visits, parents' number of drug prescriptions.

---

<sup>8</sup>The variables not shown are binary indicators for gender, education field (25 categories), drugs taken (14 categories), diagnoses (21 categories).

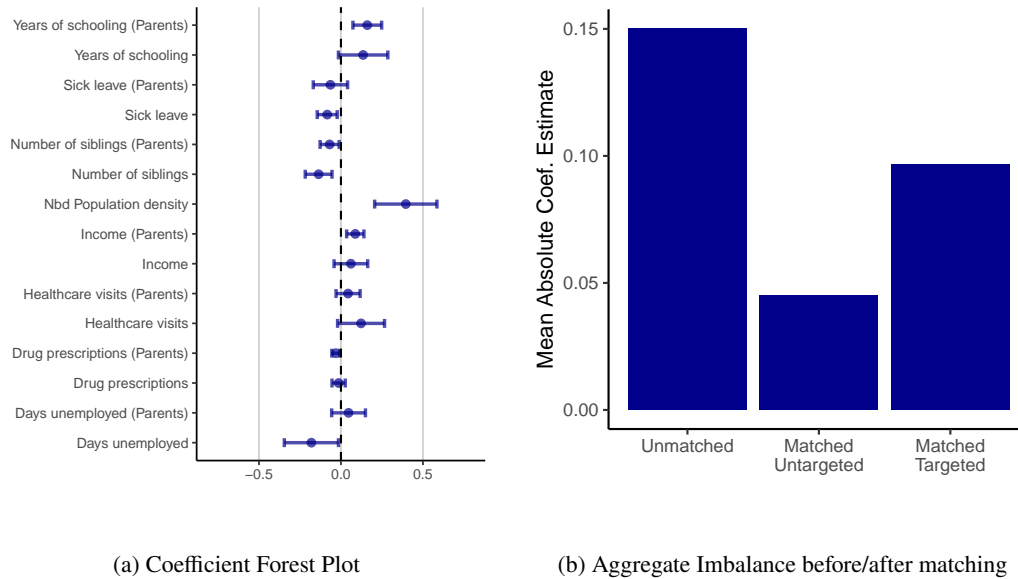


Figure 3: Pre-Treatment Covariate Balance

*Notes:* **Panel (a)** Differences in a selection of pre-determined characteristics between treated and Pandemrix control individuals. Each estimate comes from a univariate regression of the standardized covariate on the treatment status and on birth year fixed effects. Own socioeconomic characteristics, as opposed to parental characteristics, are restricted to individuals born before 1990. **Panel (b)** Mean absolute value of coefficient for all covariates used. The first column displays the mean value before matching. The second column displays the mean value for variables not selected for matching. The third column displays the mean value for variables selected for matching.

In Figure 3, Panel (b) we provide a simple statistic that summarizes the performance of the logit matching method, namely the mean absolute value of the standardized coefficients across univariate regressions from the full set of covariates (including those not displayed in Panel (a)). Overall, the imbalance decreases but, somewhat surprisingly, it decreases less for the variables used for matching (targeted) than for those that are not (untargeted).

In Figure B6 and Figure B7 we generalize Figure 3, Panel (b) by considering covariate balance across networks of individuals, separately for the pre- and post-swine flu samples. The methods include nearest-neighbor matching based on logit propensity scores (NN), coarsened exact matching (CEM), and nearest-neighbor matching using gradient boosting (XGBM). A few patterns are worth noting. First, nearest neighbor matching based on logit propensity scores generally outperforms XGBM and in particular CEM. Second, larger networks exhibit less imbalance even before matching, likely because they are more weakly

connected to individuals belonging to the group with some degree of selection into diagnosis. Third, matching achieves better balance for these larger networks as well, probably because the pool of potential matches is larger.

For the main sample, we use a set of candidate matching variables consisting of characteristics of individuals and their parents as measured between 2005 and 2009, right before individuals may have developed narcolepsy symptoms due to Pandemrix. For individuals that are born after 1990 we restrict the matching variables to demographic and parental characteristics as we cannot measure socioeconomic characteristics for the individuals themselves between 2005 and 2009, as they are too young to observe e.g., income and years of schooling. Out of the treated individuals, 278 are born 1990 or later, 77 are born before 1990. We also use parental and time-invariant characteristics for the pre–swine flu sample, as socioeconomic characteristics are likely influenced by the presence of narcolepsy (recall that we do not observe when these individuals first developed narcolepsy). The set of candidate covariates is listed and defined in Section C. We use a simple LASSO algorithm with cross-validation to select variables to match on. We use this procedure separately for individuals born 1990 or later and those born before 1990, for the pre–swine flu sample and post–swine flu sample, as well as for the different networks, implying that the variables used for matching varies across these dimensions. Matching is always done on the characteristics of the network members rather than the characteristics of the focal individuals. Note, that although we ensure that the focal individuals have received Pandemrix, we make no such restrictions on the network members.

We follow [Abadie and Spiess \(2022\)](#) in clustering standard errors at the match level, and additionally cluster at the level of the focal members in the network.

## 6 Quantifying the Aggregate Role of Healthcare Experiences

The conceptual framework is motivated by the idea that personal experiences influence vaccination decisions. Before studying the effect of exposure to narcolepsy, we provide large-scale descriptive evidence on the role of prior healthcare experiences during the COVID-19 pandemic. In Figure 4 we display AUC values for predicting whether or not individuals take at least one COVID-19 vaccine dose. We restrict the sample to individuals aged 40 to 60 in 2021 and who, based on their diagnosis history, were not in a COVID-19 risk group. This restriction implies that there is no clear medical reason for these individuals to refrain from taking the COVID-19 vaccine. We predict their COVID-19 vaccination status based on all health related variables that we observe, namely their diagnoses, medical drugs and reported side effects between 2010 and 2020. Our

preferred method for computing predicted probabilities is a recurrent neural network that explicitly captures the ordering, but not the timing, of events. We also use XGBoost but consistently get somewhat smaller AUC values. We benchmark the prediction score against what we get using (time-invariant) socioeconomic characteristics. The estimation sample contains 2.84 million individuals (90% train / 10% test).

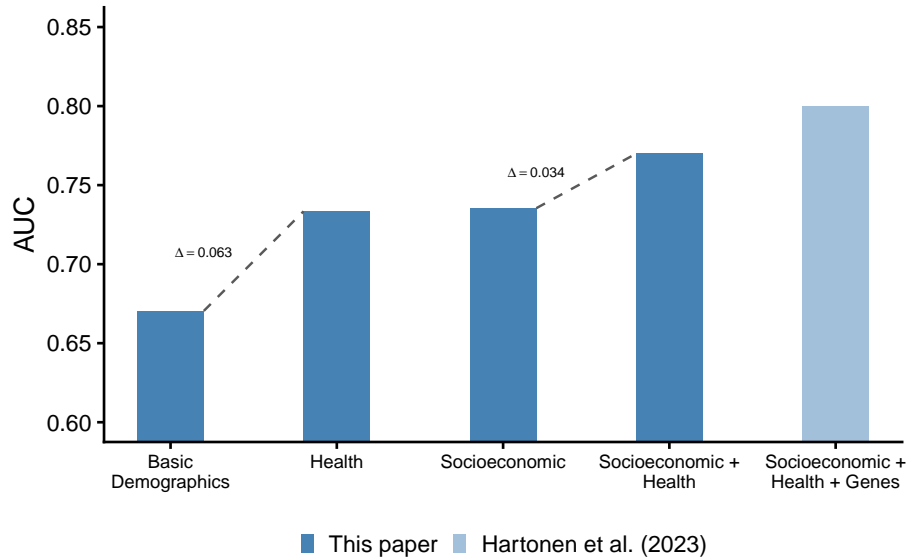


Figure 4: Predicting Who Takes the COVID-19 Vaccine

*Notes:* Out-of-sample AUC values for predicted probabilities using a recurrent neural network. The sample is restricted to individuals aged 40–60 who do not belong to COVID-19 risk groups. *Demographic characteristics* include number of siblings, gender, birth year, and geographic origin. *Socioeconomic characteristics* include income, years of schooling, days unemployed, and days sick, and are specified in addition to the demographic variables. *Health characteristics* include history of diagnoses, drug use, and reported adverse events, and likewise build on the demographic variables. All models are evaluated out-of-sample. The model used is described in detail in Section F.

We find that health history alone yields an AUC of 0.65, Socioeconomic variables yields 0.72 and the full model, including both health and socioeconomic variables, yields a value of 0.76. These numbers are similar to what is found in a recent study using Finnish data [Hartonen et al. \(2023\)](#) who find an AUC of 0.8, using data similar to ours but where the authors also include information on genes and vaccination status of family members. Similar to us, they find that, apart from income, medication history has the highest predictive power. The prediction exercise provides descriptive evidence that previous experiences with healthcare at large shape attitudes toward vaccines. Taken together, while socioeconomic characteristics strongly predict vaccination

uptake, the incremental increase from including personal health experiences—as measured by the history of diagnoses, drugs and reported adverse events—is substantive, lending empirical support to the importance of personal experiences.

## 7 Exposure to narcolepsy

### 7.1 Main Results

**Effects among diseased individuals and family members** We document the effects on our main measure of COVID-19 vaccine hesitancy in Figure 5. These measures include: (i) whether an individual takes the COVID-19 vaccine, (ii) the number of doses conditional on taking one dose, and (iii) the time until the first vaccine dose. We present results both for individuals who themselves developed narcolepsy and for their family members. Specifically, we report results for  $\tau_{post}$ , the mean difference between treated individuals (post—swine flu narcolepsy) and matched individuals that received Pandemrix;  $\tau_{pre}$ , the mean difference between placebo control individuals (pre—swine flu narcolepsy) and matched individuals from the general Swedish population; and  $\tau_{dd}$ , the difference between  $\tau_{post}$  and  $\tau_{pre}$ .

For the post—swine flu sample, used to estimate  $\tau_{post}$ , individuals who developed narcolepsy are 40 percentage points less likely to have received at least one vaccine dose than their matched controls, among whom 92 percent are vaccinated. We also find a lower vaccination rate for the placebo control group of individuals that developed narcolepsy before the Pandemrix vaccination campaign. As discussed in Section 5, it is unlikely that there something specific to having narcolepsy that cause these individuals to abstain from vaccination. Instead, the plausible explanation is that the Pandemrix-narcolepsy episode is more salient to these individuals and that they therefore are more hesitant. Combining  $\tau_{post}$  and  $\tau_{pre}$  yields a large effect of 35 percentage points lower vaccination rate. Given the informational spillovers, we interpret this estimate as a lower bound on the effect of developing Pandemrix-induced narcolepsy.

To gauge the economic significance of our estimates, we compare the magnitude of the effects to the socioeconomic gradient in vaccination rates which is shown in Figure B1. While there is a sizable gradient in immunization outcomes along socioeconomic characteristics, it is small compared to the effect sizes that we find. For example, the difference in vaccination rate between individuals with post-tertiary schooling and individuals with only primary school education is no more than five percentage points.

We next consider the number of doses, conditional on receiving at least one dose, to account for vaccine

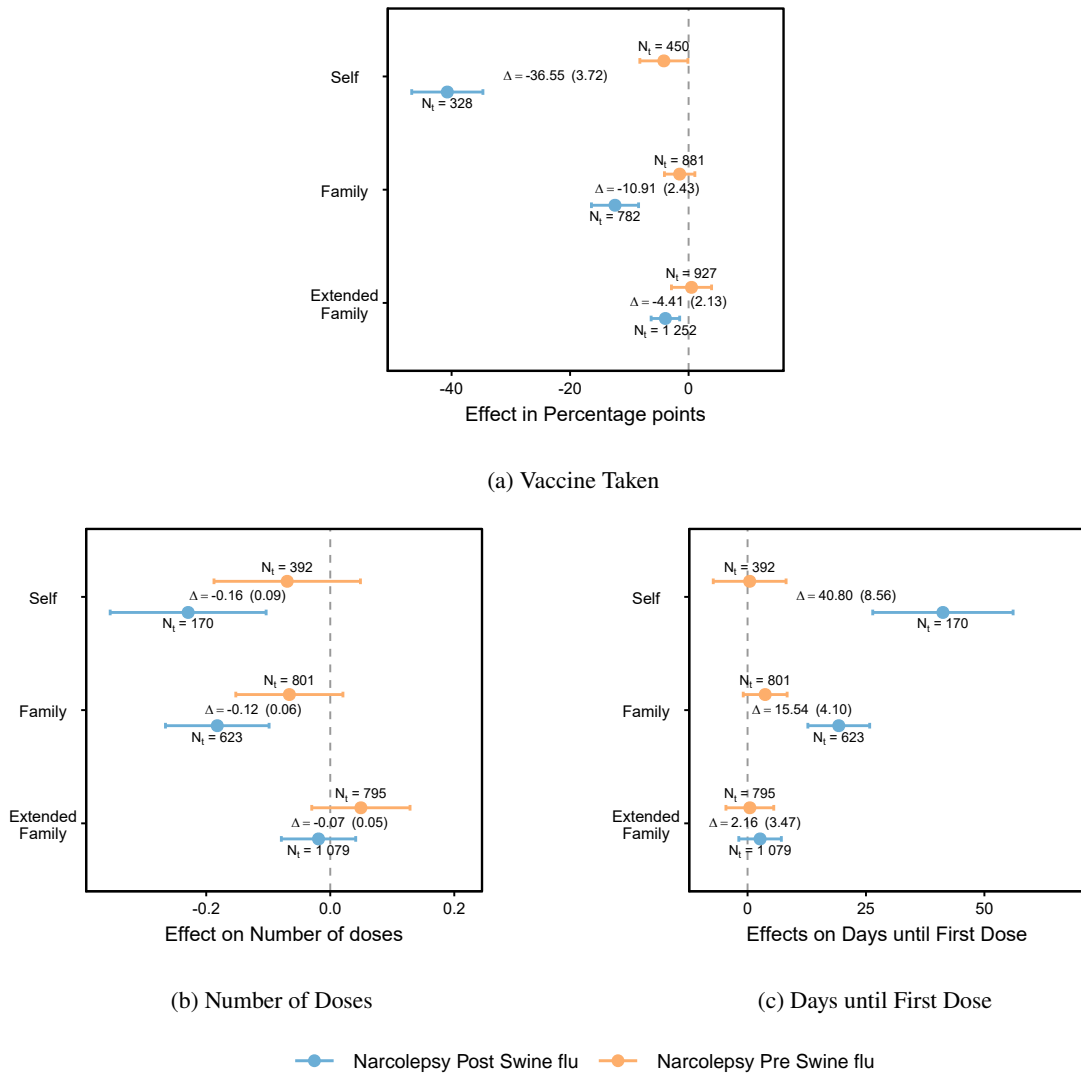


Figure 5: Main Results—Diseased Individuals and Family Members.

Notes: This figure displays coefficients corresponding to  $\tau_{post}$  and  $\tau_{pre}$  in eq. (4). First row corresponds to individuals that themselves developed narcolepsy (Self). The second and third rows show the corresponding coefficients for Family members, defined as siblings and parents, and Extended family members, defined as aunts/uncles and cousins. **Panel (a):** Ever taken at least one COVID-19 dose. **Panel (b):** Number of doses, conditional on at least one dose. **Panel (c):** Days elapsed until the first COVID-19 dose. Standard errors are clustered by treatment cluster, defined by the treated or control individual to whom a network member is related as well as at the match level.

hesitancy expressed in incomplete vaccination schedules (the intensive margin). We find differences of 0.1–0.4 (2.8 doses in the matched control group of diseased individuals) fewer doses. These coefficients are small in magnitude compared to the effects on the extensive margin. In other words, individuals willing to take the first dose appear do not appear much more hesitant to complete the full vaccination schedule.

Turning to the timing of vaccination, diseased individuals delay their first vaccination by roughly 40 days, family members by around 10 days, and extended family members by about three days. The fact that affected individuals delay their vaccination is an indication that they are concerned about adverse events and prefer to wait and see if others experience adverse events from the COVID-19 vaccine. The effects for days elapsed for close family members is comparable in magnitude to the gap between individuals with only primary schooling (nine years) and those with post-tertiary education.

We replicate the results in Figure 5 using coarsened exact matching and propensity scores estimated via gradient boosting in Figure E1 and Figure E2. The findings remain virtually identical with the exception of *Number of Doses* where we find negligible effects for CEM. In Figure B9 we address the concern that placebo control individuals are significantly older than treated individuals. This is a concern if the mediating effect of *having narcolepsy* differs by age. We display one estimate where we reweight the birth year distribution of narcolepsy individuals in the placebo control sample to match the birth year distribution of the treated units in the main sample. To account for the limited overlap in age between the two groups, we fit a spline of vaccination uptake on age using individuals in the pre–swine flu placebo sample. We then use this model to predict vaccination uptake for the age distribution observed in the post–swine flu treatment sample. The resulting counterfactual mean is compared to the observed placebo mean to assess whether the estimated effect,  $\tau_{pre}$ , changes. The results are nearly identical, suggesting that our estimate of  $\tau_{pre}$  indeed captures the counterfactual effect of developing narcolepsy—unrelated to Pandemrix—at the age when individuals developed Pandemrix-induced narcolepsy.

In light of our framework, the difference in magnitude of the results between direct family members and extended family members is puzzling. Both direct and extended family members are regularly exposed to the individual with narcolepsy, and know that it may have been induced by Pandemrix. Direct family members, however, are more likely to witness the patient’s sudden change in lifestyle, ongoing symptoms and daily challenges, which makes the condition much more salient for them. This is particularly true, because direct family members bear the caregiving responsibilities of the diseased, such that their memories of these struggles become more vivid and influential in shaping subsequent health decisions. We further hypothesize that the heightened emotional tagging among close relatives of the patients makes the experience particularly

easy for them to recall and use for simulation.

**Learning about predisposition** The finding that direct family members have similar but weaker effects on immunization outcomes as illustrated in Figure 5 may be partially explained by close family members sharing genetic traits with diseased individuals, leading them to perceive that they possess predispositions that make them more susceptible to adverse events from vaccination. In Figure B11, we show results for the partners of the diseased individuals. For partners, we find large estimates within the same range of those found for close family members of diseased individuals. Keeping in mind the endogeneity in partner choice, we interpret these findings as evidence that the results in Figure 5 for network members are primarily driven by information about the risk of severe adverse events, rather than by learning about one's predisposition to adverse events.

In contrast to the results for partners, Figure B12 shows how the effect on COVID-19 vaccination varies by year of narcolepsy onset. To ensure comparability with the pre-swine flu period, we do not restrict the sample to individuals vaccinated with Pandemrix. The effect fades over time and is largest among those who developed narcolepsy soon after the swine-flu pandemic. Individuals who developed narcolepsy later on were likely aware of the Pandemrix-narcolepsy episode but appear less hesitant, suggesting that they viewed their condition as less likely to be vaccine-induced and therefore saw themselves as less likely to have a predisposition to adverse events.

**Mechanisms alternative to altered perceived costs** Our simple conceptual framework suggests that treated individuals respond primarily to altered beliefs about the risk of adverse events, although an alternative explanation consistent with the previous finding is that treatment instead changes perceived infection risk, and thus beliefs about the benefits of vaccination. Lower perceived benefits would, for example, be consistent with individuals shying away from the healthcare system altogether. We elicit the perceived benefits of vaccination through two measures: (i) the number of COVID-19 self tests that they take, and (ii) whether they make phone calls for medical advice related to COVID-19 during the pandemic. Because of differences in practices across regions regarding both testing and medical advice, we match treated individuals to individuals living in the same municipality. In doing so, we select control individuals from the overall Swedish population rather than from those who received Pandemrix.

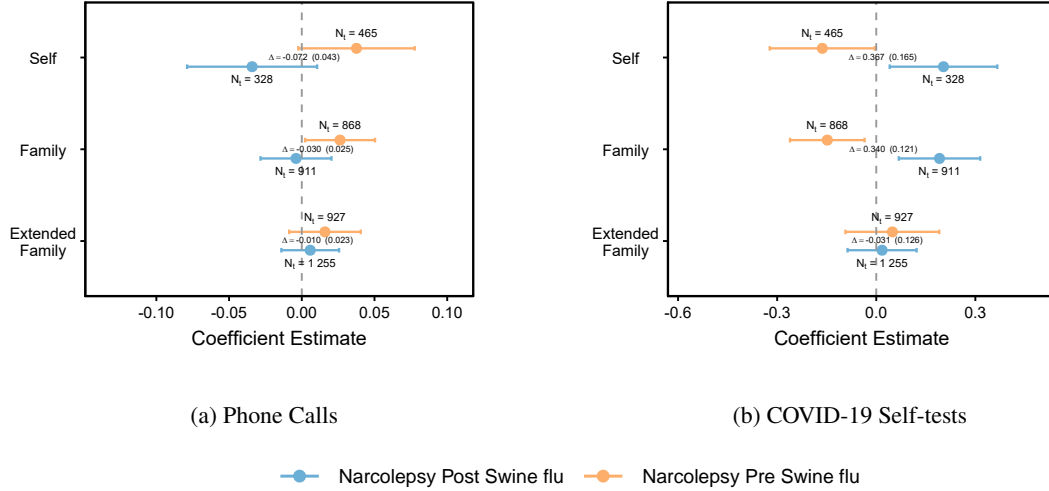


Figure 6: Eliciting the Perceived Benefits of COVID-19 vaccination

Notes: This figure displays coefficients corresponding to  $\tau_{post}$  and  $\tau_{pre}$  in eq. (4). The first row corresponds to individuals who themselves developed narcolepsy (Self), while the second and third rows show the corresponding coefficients for Family members and Extended family members, respectively. **Panel (a)** refers to whether an individual made a phone call for medical advice regarding COVID-19 symptoms during the pandemic, and **Panel (b)** to the number of COVID-19 self-tests an individual has taken.

The results are displayed in Figure 6. If anything, diseased individuals are less concerned about developing COVID-19 both measured by the number of tests and number of phone calls. The coefficients are small in magnitude but suggest that the main results may partly reflect differing assessments of the risk of contracting COVID-19 or developing severe symptoms from COVID-19, rather than solely by the perceived risk of adverse events. A simple interpretation consistent with these findings is that treated individuals may limit social exposure, thereby reducing their risk of contracting COVID-19.

**Experiencing the benefits of vaccination** We extend the analysis by examining individuals who experience the benefits of vaccination rather than its costs. Specifically, we consider individuals who develop diseases that vary in their similarity to COVID-19. The idea is that those who contract potentially vaccine-preventable conditions become more aware of the consequences of foregoing vaccination. We focus on individuals who develop conditions with symptoms similar to COVID-19 that are either (i) vaccine-preventable or (ii) infectious but not vaccine-preventable. In Figure B13, we test whether prior exposure to the benefits of vaccination affects subsequent vaccine uptake. We use the same outcome variables as before but redefine the treatment variable

to take the value 1 if an individual was hospitalized between 2005 and 2021—just before the start of the COVID-19 vaccination campaign—with one of several vaccine-preventable or related conditions: COVID-19, influenza, upper respiratory tract infection (URTI), pulmonary embolism, or sinusitis. Controls are nearest neighbors (estimated using gradient boosting) matched on socioeconomic and health characteristics, with exact matching on gender and birth year. Each regression is reweighted to match a common age distribution, ensuring that differences in coefficients across conditions are not driven by age. Our results show that exposure to the benefits of vaccination reduces vaccine hesitancy. In particular, individuals who developed COVID-19 during the initial stages of the pandemic exhibit significantly higher vaccination rates. Qualitatively similar but smaller effects are found for influenza—also vaccine-preventable and infectious—and for upper respiratory tract infections, which are often infectious but rarely vaccine-preventable.

**Similarity-based learning** Up until now we have provided little room for behavioral explanations, different from the cost-benefit analysis, for the effects that we find. One such behavioral mechanism is that individuals identify more strongly with the affected person if they are similar, and thus can easier simulate the event of experiencing an adverse event from the COVID-19 vaccine. We explore this hypothesis in Figure B15 by comparing cousins who share the same gender and age as the individual with narcolepsy to those cousins who do not. We find no evidence of larger effects among those who are more similar.

**Effects beyond family members** We proceed by providing results for extended, larger, networks. The goal of this exercise is to gain suggestive insights into the large-scale impact of the narcolepsy health scandal. If personal networks are permeable to the effects of exposure, extending beyond close family, this may indicate that aggregate vaccination uptake is likely to be lower in the wake of adverse events scandals. We construct networks so that the network members are plausibly aware of the individual who developed narcolepsy. Since most of the individuals were in school when they developed narcolepsy, we focus on schoolmates and neighbors in 2011, during the time the scandal unfolded and most of the individuals started experiencing symptoms. For colleagues, we instead restrict our attention to individuals that were colleagues (working at the same plant) around the time of COVID-19 vaccination, in 2021. These networks are moderately large, where treated individuals have on average 300 schoolmates, 33 colleagues, and 298 neighbors. Figure B10 show results for colleagues, schoolmates, and neighbors. We find no indication that vaccine hesitancy spreads within these extended networks, as measured by either of our indicators of COVID-19 vaccine hesitancy. Pandemrix-induced narcolepsy should be salient for these individuals when deciding whether to take the

COVID-19 vaccine, yet they appear unaffected. Because interaction is much more frequent within families than in these broader networks, we interpret the absence of effects here as evidence that frequency of exposure fundamentally shapes the salience and ease of recall of the narcolepsy episode, and thereby its influence on subsequent vaccination decisions.

**Attributing other diseases to Pandemrix** As stated in the data section, the best available evidence suggests that Pandemrix caused no diseases other than narcolepsy. Nonetheless, individuals may still attribute other diseases to Pandemrix, particularly if there were prior concerns linking the vaccine to that disease. We examine this by considering a number of diseases that were studied soon after the swine flu vaccination campaign (Persson et al., 2014). In Figure B16 we display COVID-19 vaccination rates by the year in which it was first diagnosed for each condition. In short, we find no indication that individuals attributed their disease status to Pandemrix, once again highlighting that individuals are rational in the sense that they only react to adverse events that were confirmed to be associated with narcolepsy.

**Healthcare visits** Finally, we examine if treated individuals refrain from the healthcare system altogether. In Figure B17 we display the yearly number of specialist healthcare visits across time. In a short period after developing narcolepsy, patients make more visits to medical professionals, likely in the examination phase, before they are diagnosed with narcolepsy. After that, the number of non-narcolepsy related healthcare visits reaches a level comparable to the control group. These results provide evidence that, although individuals experience a very severe adverse event in the form of Pandemrix-induced narcolepsy, they do not shy away from the healthcare system altogether.

**Taking stock** Affected individuals and their family members display a large difference in COVID-19 immunization outcomes. We find similarly large effects among partners of diseased individuals, suggesting that individuals learning about a potential pre-disposition to adverse events is not a key driver of the negative effects on immunization outcomes. We find little evidence of the theoretical possibility that results are explained by differing assessment of the benefits of vaccination. Past exposure to the benefits of vaccination as measured through individuals that experience influenza, rather than the costs, leads to improved immunization outcomes. We find no effects on networks beyond family members, suggesting that the *frequency* of exposure to the Pandemrix–narcolepsy episode is key. Yet, the broad take-away is that individuals place substantial weight on their own experiences when confronted with novel risks. A central question is why. Graeber et al. (2024) provides evidence of stark differences in the ability to recall information that are conveyed through

stories as opposed to statistics; stories are more vivid and easier to recall. In our setting, the exposure to narcolepsy may impede the ability to recall relevant statistics about risks of adverse events.

## 7.2 Heterogeneity in Size of Database & Health Literacy

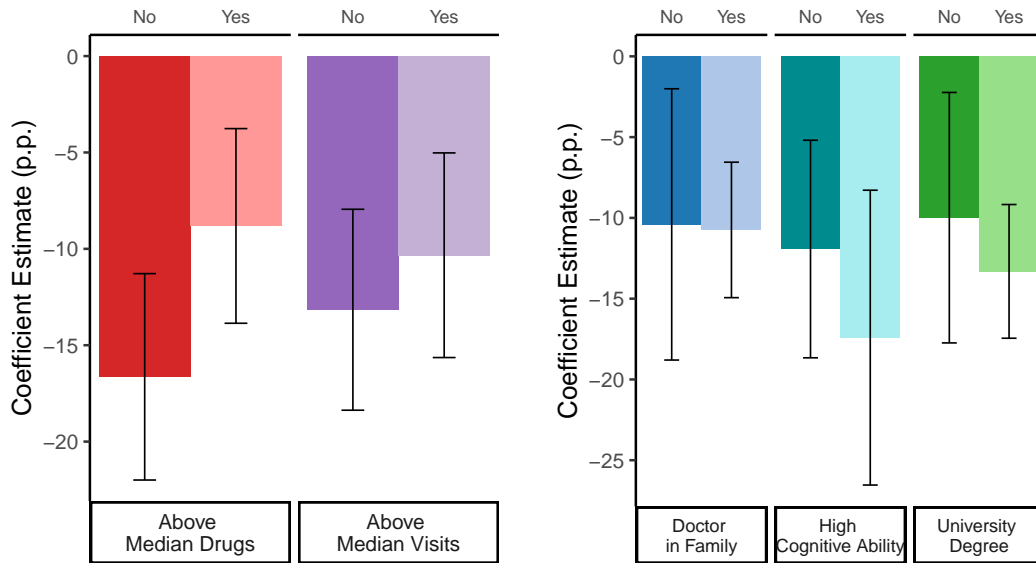
We now test two implications of our conceptual framework as laid out in Section 3. Prediction 1 states that the immunization response should be stronger among individuals with low health literacy, whereas Prediction 3 posits a stronger response among those with fewer prior healthcare experiences—that is, a smaller experience database. We focus on *Vaccine Taken* for family members to individuals that develop narcolepsy as to maximize statistical power.

In Figure 7, Panel A we examine the role of the size of the database. We consider experiences in the healthcare domain under the assumption that these are the relevant experiences that could both influence an individual’s perceived risk of COVID-19 vaccine but also interfere with recalling experiences of adverse events. We use two proxies for experiences in the healthcare domain: (i) The number of unique drugs taken, (ii) The number of healthcare visits.

We find that Individuals who have taken a greater number of unique drugs react somewhat less to exposure to narcolepsy. We observe no such differences for number of healthcare visits. If anything, this supports the notion that other relevant experiences interfere with the experience of having a family member with narcolepsy. In Figure 7, Panel B we examine the role of health literacy which in turn governs how much individuals rely on official information ( $\theta$  in our model). We use three different proxies for health literacy: whether individuals (i) have a doctor in the family, (ii) have above-median cognitive ability, and (iii) have a university degree. Taken together, the estimates indicate that higher health literacy does not dampen the behavioral response: better-informed individuals are no less likely to reduce their immunization uptake after an adverse event.

What may explain the lack of smaller effects for individuals with high health literacy? A possible counteracting force that would imply *larger* effects for individuals with high health literacy goes as follows: Individuals without experiences—good or bad—are more receptive to information about vaccine safety, and high health literacy makes information easier to access. This would imply a positive association between health literacy and vaccination (as documented in Figure B1). On the other hand, among individuals with negative experiences, personal experience crowds out official information, so perceived risk no longer varies with health literacy. We think that this counteracting force is the most plausible explanation for the results in Figure 7, Panel B.

Furthermore, the lack of heterogeneity with respect to health literacy are reminiscent of [Malmendier](#)



(a) Size of Database

(b) Health Literacy

Figure 7: Testing Predictions From Model: Health Literacy & Size of Database.

Notes: This figure shows results for two of the predictions from the model: **Panel (a)** The role of previous healthcare experiences and **Panel (b)** The role of health literacy, which in our model manifests itself through  $\theta$ . We restrict the attention to  $\tau_{post}$  for family members and for the binary outcome variable *Vaccine Taken*. *Above Median Drugs* is defined as "Yes" if an individual has above median number of drugs taken ( $\approx 6$ ) between 2005 and 2009 relative to its birth year peers. *Above Median Visits* is defined equivalently but for visits to specialized healthcare between 2005 and 2009 ( $\approx 0.6$ ). *Doctor in family* is defined as "Yes" if an individual has a parent or a sibling with a medical degree or a nursing degree. *High Cognitive Ability* is based on cognitive tests completed by military conscripts. It is equal to "Yes" if an individual had above median score on the test compared to peers born the same year. *University Degree* is equal to "Yes" if an individual has at least a bachelor degree, corresponding to three years of higher education.

et al. (2021) who shows that central bankers, arguably a group with expertise in inflation forecasting, let their own inflation experiences shape their forecasts. Archibong and Annan (2023); Anderberg et al. (2011); Chang (2018) likewise find larger effects on vaccine hesitancy among individuals with higher education after exposure to medical scandals.

## 8 Exposure to General Severe Adverse Events

Up to this point, we have focused on the effects of exposure to narcolepsy. This was a particularly severe adverse event, and the deployment of the swine flu vaccine shared notable similarities with the deployment of the COVID-19 vaccine. Moreover, we consider it highly likely that developing and reporting narcolepsy as an adverse event is orthogonal to factors influencing vaccine hesitancy, as narcolepsy has few comorbidities and individuals had strong incentives to report it as an adverse event. We now broaden the analysis to consider the effects of severe adverse events from all licensed pharmaceuticals (including all vaccines apart from Pandemrix).

Unlike the rare episode of Pandemrix-induced narcolepsy, the focus on serious adverse events that recur across drugs, countries, and time allows us to address the policy-relevant question of how routinely experienced adverse events shape future vaccine hesitancy. Because Swedish healthcare professionals are expected to report all suspected adverse events, regardless of whether the symptom is already recognized as vaccine-related, the register of suspected adverse events captures a broad and plausibly representative spectrum of adverse reactions. By focusing on events classified as serious and reported by physicians, our analysis addresses reactions often severe enough to necessitate hospitalization, those most likely to shape subsequent immunization decisions.

We focus on 1 700 individuals for whom vaccine adverse events were reported and 33 000 individuals for whom non-vaccine adverse events were reported between 2014 and 2020, just before COVID-19 vaccination began in Sweden. As displayed in Figure 8, individuals developing and reporting adverse events differ from individuals with the same age and gender in a variety of ways. For the sake of simplicity, we focus on a handful of health and socioeconomic variables. Individuals reporting adverse events appear to be generally sicker, as they exhibit more sick leave days, healthcare visits and drug prescriptions, but do not appear positively selected on socioeconomic characteristics.<sup>9</sup> Individuals that later reported adverse events are *more* likely to have taken the Pandemrix vaccine.<sup>10</sup> This is evidence that our results are unlikely to be explained by latent

<sup>9</sup>Individuals that report mild adverse events are, however, positively selected in terms of socioeconomic characteristics.

<sup>10</sup>Pandemrix vaccination status is observed only for individuals living in the health care regions for which we have data access, which

vaccine hesitancy that correlates with both the propensity to report adverse events and willingness to take the COVID-19 vaccine.<sup>11</sup>




Figure 8: Balance in Pre-treatment Characteristics—General Adverse Events

*Notes:* This table displays balance pre-treatment characteristics among individuals reporting general (vaccine and non-vaccine) adverse events before and after matching on propensity scores where the propensity scores are computed using a neural network. Each coefficient estimate comes from a univariate regression of the standardized covariate on the treatment status and on birth year fixed effects. Pandemrix taken is defined for a subset of individuals who lived in healthcare regions in 2009 where we access individual level data on Pandemrix vaccination. *Pandemrix taken* is deliberately excluded from the computation of propensity scores as to display how well balance in immunization outcomes during the swine flu pandemic can be achieved using other health- and socioeconomic characteristics.

We identify a control group of comparable individuals in two steps. First, we find candidate control individuals who took the same drug in the same year and who share the same birth year. Within these *birth year* × *drug* × *drug year*-cells, we then perform nearest-neighbor matching, with propensity scores computed using a neural network based on socioeconomic and health variables.

The individuals who report general adverse events are older than those who developed Pandemrix-induced narcolepsy. In 2021, when the vaccine was rolled out, the median age among reporters was 63. At this age, forgoing or delaying COVID-19 vaccination entails substantial health risks relative to the individuals who developed narcolepsy.

allows us to assess balance but prevents us from using it for matching in the full sample.

<sup>11</sup>The positive relationship between reporting adverse events and taking Pandemrix is unchanged when we control for simple health characteristics such as number of drugs, numbers of healthcare visits and number of days on sick leave.

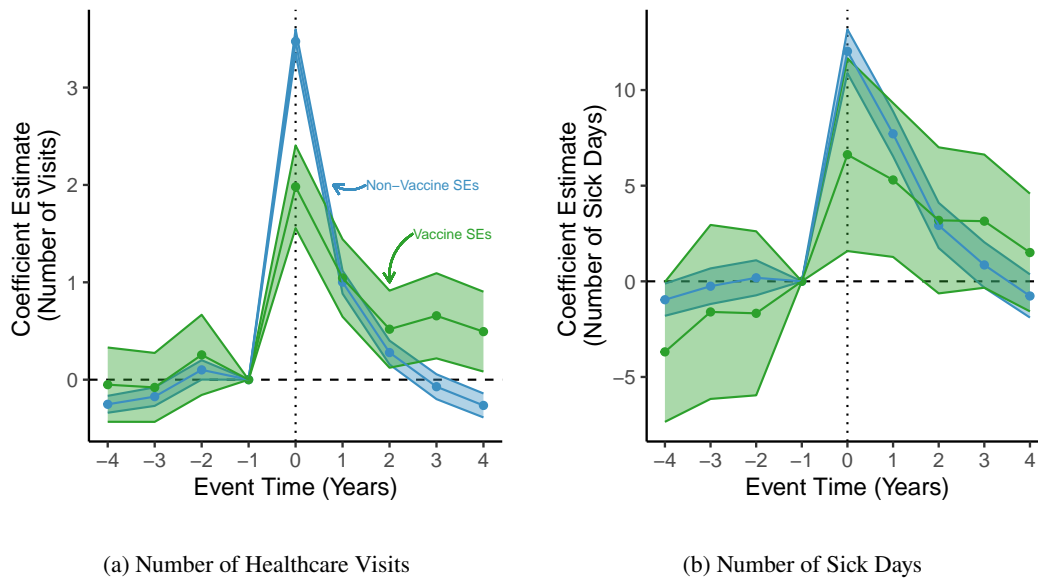


Figure 9: DiD-Estimates of Health Effects from Adverse Events

*Notes:* The figure shows estimated health outcomes around the reporting of a severe adverse event relative to a matched control group split up by vaccine adverse events (green) and non-vaccine adverse events (blue). **Panel (a)** reports the number of visits to specialized healthcare; **panel (b)** reports the number of sick days. Estimates are obtained from a standard two-way fixed effect specification with individual level and calendar year fixed effects.

What does experiencing an adverse event entail? In Figure 9 we provide results from a simple difference-in-difference strategy showing the effects on health outcomes around the time an adverse event is reported, relative to a matched control group. Experiencing an adverse event is associated with worsened health outcomes along observable measures: on average, affected individuals have three additional visits to specialized healthcare and twelve more days of sick leave in the year of experiencing an adverse event. These effects also extend up to two years after the event year. Importantly, when we restrict the attention to adverse events from vaccines, the effects on health are roughly the same. This implies that differences in effects on vaccine hesitancy between vaccine and non-vaccine adverse events are unlikely to stem from differences in severity.

We now turn to the results for non-vaccine adverse events and vaccine-related adverse events, displayed in Table 1. The treatment variable equals 1 if a suspected adverse event for a vaccine or another drug was reported for an individual between 2014 and 2020, and 0 otherwise.

Table 1: Adverse Events and COVID-19 Immunization Outcomes

	Non-Vaccine Adverse Events			Vaccine Adverse Events		
	Vaccine Taken	Number of Doses	Days Elapsed	Vaccine Taken	Number of Doses	Days Elapsed
<b>Diseased Individuals</b>						
adverse event	-0.006 (0.003)	-0.038 (0.015)	-3.19 (0.68)	-0.075 (0.019)	-0.056 (0.073)	10.3 (3.51)
N. Treated	12 230	11 326	11 326	672	553	553
<b>Family Members</b>						
adverse event	0.002 (0.002)	0.030 (0.010)	0.30 (0.42)	-0.037 (0.011)	-0.092 (0.040)	5.91 (1.92)
N. Treated	28 509	26 323	26 323	1 643	1 431	1 431
<b>Children</b>						
adverse event	0.009 (0.004)	0.021 (0.009)	-0.74 (0.56)	-0.042 (0.022)	0.056 (0.050)	7.42 (3.55)
N. Treated	24 188	19 670	19 670	669	513	513
<b>Partner</b>						
adverse event	-0.001 (0.003)	-0.008 (0.020)	-1.32 (0.79)	-0.009 (0.015)	-0.056 (0.096)	7.90 (4.14)
N. Treated	6 216	5 951	5 951	245	237	237

*Notes:* This table displays results from regressing the COVID-19 immunization outcome variables on a variable indicating if someone developed an reported an adverse event between 2015 and 2020. Columns 1–3 use an indicator for reporting any adverse event between 2015 and 2020 as the regressor (“Non-vaccine Adverse Events”). Columns 4–6 instead use an indicator for reporting a vaccine adverse event (“Vaccine Adverse Events”). Each treated individual is matched 1:1 to an untreated individual who (a) was born in the same year and (b) took the same drug in the same year as the treated individual. within the exact-match cells, treated and control units are further paired by nearest-neighbor matching on a propensity score built from pre-treatment socioeconomic and health variables.

For diseased individuals, the broad pattern is that these individuals become more hesitant with respect to

the COVID-19 vaccine as measured by our three immunization outcomes. The coefficients for vaccine adverse events are roughly proportional to those found for narcolepsy with the coefficients for vaccine taken and for days elapsed being about one-third as large as the corresponding coefficients for narcolepsy. For family members, we find precisely estimated effects close to zero for non-vaccine adverse events, but meaningful effects for vaccine adverse events. Recall that there is no meaningful difference in severity between non-vaccine adverse events and vaccine adverse events. Instead, the larger effects for vaccines stem from the perceived similarity to COVID-19 vaccines which boosts both recall and simulation compared to non-vaccine adverse events. One concern is that individuals reporting vaccine adverse events are, on average, younger (mean age = 39, SD = 25) than those reporting other adverse events (mean age = 55, SD = 20). Since younger individuals face lower health risks from remaining unvaccinated, this age difference may affect the comparability of the groups. In Table B2 we reweight observations in the vaccine sample to match the age distribution of individuals in the non-vaccine adverse events-sample. The coefficients decrease only marginally from doing so.

**Stickiness of experiences** In Figure B14, we break down the results by the year the adverse event was reported, between 2013 and 2020. We reweight observations to match a common age and drug distribution so that differences in estimated effects are not driven by changes in the sample composition over time. There is no indication that the effect decays, suggesting that although adverse events are, perhaps, not on top of mind, their memories are reactivated years later when individuals face new vaccination decisions. A concern is that the salience of the narcolepsy episode interacts with the timing of adverse events—that is, individuals who developed adverse events soon after the narcolepsy episode unfolded may have been more affected. This interaction could explain the dip in uptake among those who experienced vaccine adverse events in 2016, when public debate over compensation resurfaced as the government assumed responsibility for affected individuals.

**Learning about predisposition** If individuals learn about a personal predisposition to adverse events after experiencing one, abstaining from future vaccines can be rational. While we do not take a stance on whether abstention is rational in our setting, we assess whether individuals behave as if learning about a predisposition by comparing adverse events that are *type-learning*—the individual learns she is generally prone to adverse events—with those that are *idiosyncratic*—the event is likely limited to that occasion or drug. Because it is difficult to determine, for each drug–adverse event combination, the probability of an adverse event from the COVID-19 vaccine, we classify drug–adverse event combinations using a large language model. Details about

the classification and examples of type-learning and idiosyncratic adverse events are provided in Section H. This approach scales to many combinations and broadly reflects expert knowledge—and thus the affected individual’s perception—of whether a combination is idiosyncratic. We construct a binary indicator that splits treated individuals 50/50 across the two classes, maximizing power while remaining agnostic about the threshold between idiosyncratic and type-learning. In regressions of subsequent vaccination, the effect is -0.147 (SE 0.028) for type-learning adverse events and -0.0679 (SE 0.023) for idiosyncratic adverse events; the difference is 0.079 (SE 0.036),  $p \approx 0.03$ . These results suggest that individuals react more strongly to adverse events that likely reveal a predisposition.

## 9 Spillover to Children’s Vaccines

We proceed to examine the effects of exposure to adverse events on vaccines against diseases other than COVID-19. We focus on children’s vaccines against measles, mumps and rubella (MMR), diphtheria and pneumococcal disease.<sup>12</sup> The purpose of this exercise is twofold.

First, and most importantly, we assess whether the effects we observe for the COVID-19 vaccine are specific to newly introduced pharmaceutical products. Because external information is sparse in unprecedented situations such as COVID-19, individuals rely heavily on their own prior experiences. The same mechanism is less likely to apply to established drugs, such as common child vaccines, which have historically low and well-documented incidences of severe adverse events.<sup>13</sup> As a consequence, access to scientific evidence is readily available and information costs are low compared to the case of COVID-19 vaccines. Against this background, we interpret any change in vaccination outcomes from adverse events related to established child vaccines as suggestive evidence of a shift in the degree to which individuals rely on their own experience  $\theta$  vis-à-vis best available evidence.

Second, we explore whether the effect of adverse events on vaccine outcomes translate into situations where health decisions are made on behalf of others and individuals do not bear the immediate consequences of their decision. Parents have a mandate over health care decision of their children from an early age and thus play a crucial role on whether the children ever come in contact with the public healthcare system throughout

---

<sup>12</sup>The measles vaccine is almost always given in combination with the mumps vaccine and the rubella vaccine, combined called the MMR vaccine. We focus on measles vaccinations, but it will typically coincide with vaccination against also mumps and rubella.

<sup>13</sup>The diphtheria vaccine has been included in the program since the 1940s, and measles—via the MMR vaccine—since 1982, with a standalone measles vaccine available from 1972. The pneumococcal vaccine was added to the Swedish program in 2009, though available in the US since 1977.

their childhood.

We define a sample of individuals born in 2013 or later who resides in Sweden throughout their entire lives. The three standard vaccines we focus on are administered for the first time below two years of age. This includes the measles vaccine, first administered at 18 months; the diphtheria vaccine, first due at 3 months, and most often administered in combination with polio, tetanus, Hepatitis B, pertussis, and influenza type B vaccines; the pneumococcus vaccine, which is also due at 3 months of age. Both pneumococcus and diphtheria are on average first administered at approximately 3.2 (SD=1.74) and 3.7 (SD=5.73) months of age in our data set. The average first vaccination age for the measles shot is in line with administrative regulations showing an average age of 19 (SD= 12.02) months. Our adherence shares resemble official numbers by the Swedish public health authority ([Folkhälsomyndigheten, 2024](#)). A share of individuals abstain from taking the vaccines in the standard children’s vaccine program altogether. Notably, among 5.9% of children, parents abstain from having the measles vaccine administered, compared to 1.3% for diphtheria and 2.3% for pneumococcus. The higher share for measles compared to diphtheria and pneumococcus is likely a result of the perceived link between measles and autism which we return to in section 10. Panel B of Table B3 illustrates that children’s vaccination decisions correlate highly, suggesting that parental vaccine hesitancy reflects a general sentiment rather than concerns about specific vaccines.

## 9.1 Exposure to Narcolepsy

We study effects among children as well as nieces and nephews of individuals that developed narcolepsy. The dependent variable equals one if a child adheres to the vaccination schedule—defined as receiving one dose of the MMR vaccine and three doses each of the pneumococcal and diphtheria vaccines—by age two. While the data does not allow us to observe whether children receive the vaccine later in adolescence or adulthood, those born early in our sample period (2013–2014) who receive least three doses, 89% do so before age 2. The corresponding figure for pneumococcus is 99%. Together, these patterns suggest that few individuals classified as non-adhering by our measure later go on to receive the recommended number of doses. Table 2 reports results for children (Panel A) and for nieces and nephews (Panel B) of individuals who developed narcolepsy and their matched controls. Column 5 contains the estimates for an aggregate index, defined as the average vaccine adherence for the three different diseases examined. Given that the child vaccination data begins in 2013, we are restricted to a quarter of the main sample who have children born after that year. Given this limitation, we deploy a simple matching strategy where we match exactly on coarsened versions of socioeconomic characteristics of the parents as well as birth month of the child. We

focus on the first difference  $\tau_{post}$ , and therefore do not remove potential effects from being a network member of someone who developed narcolepsy prior to the swine flu pandemic. Since we regress a binary outcome variable on a binary treatment variable, we augment the OLS with odds ratios along with confidence interval computed using Fisher’s exact test. This avoids reliance on asymptotic normality for the odds-ratio. While these estimates are not statistically significantly different from zero, the negative point estimates are large across the three different vaccines, indicating that individuals who develop narcolepsy partly also abstain from fully vaccinating their children with well established vaccines.

Table 2: Exposure to Narcolepsy and Children’s Vaccine Outcomes

	Measles	Diphtheria	Pneumococcus	Index
<b>Table A: Children</b>				
OLS	-0.055 (0.038)	-0.041 (0.038)	-0.057 (0.047)	-0.051 (0.036)
Odds Ratio	0.612 [0.301, 1.409]	0.723 [0.329, 1.878]	0.652 [0.337, 1.382]	
Num. Treated.	64	64	64	64
<b>Table B: Nieces and Nephews</b>				
OLS	0.004 (0.011)	0.013 (0.008)	0.005 (0.019)	0.007 (0.008)
Odds Ratio	0.727 [0.443, 1.264]	0.768 [0.457, 1.381]	0.853 [0.541, 1.406]	
Num. Treated.	157	157	157	157

*Notes:* This table presents coefficient estimates for the effect of exposure to a relative with narcolepsy on children’s vaccine outcomes, separately for own children and for nieces and nephews. The regressions are run on a matched sample where each child is matched exactly on birth month and where parents are matched on coarsened versions of parental years of schooling, education field, income, number of healthcare visits and number of drugs prescribed. The dependent variable is equal to one if an individual adheres to the vaccination schedule. Confidence intervals computed using Fisher’s exact test in square brackets.

We interpret these results as suggestive evidence that trust in information regarding the safety of vaccines

provided by healthcare authorities and scientific authorities plays a mediating role in vaccine hesitancy (indeed this mechanism may partially explain the results for COVID-19 vaccine decisions). As illustrated in Panel B, we find no similar effects for nieces and nephews, indicating once again that the impact of vaccine-induced narcolepsy on overall immunization outcomes is limited with regard to personal networks.

## 9.2 Exposure to General Adverse Events

In line with our analysis of narcolepsy vis-à-vis other adverse events in the previous sections, we now consider the effects from other severe adverse events. Once again, we consider the universe of adverse events (excluding reports from Pandemrix and COVID-19 vaccines). Importantly, the longitudinal data structure allows us to account for unobserved time-invariant vaccine hesitancy by exploiting the variation in having a parent who have experienced an adverse event across children within the same family. To this end, we estimate the following regression equation:

$$y_{pb} = \delta_b + \gamma_p + \beta T_{pb} + \epsilon_{pb} \quad (6)$$

where  $\delta_b$  denotes birth-order fixed effects,  $\gamma_p$  denotes parent fixed effects, and  $T_{pb} = 1$  if parent  $p$ 's child of birth order  $b$  was born after the parent reported a severe adverse event, and 0 otherwise.  $y_{pb}$  denotes the child's vaccination status, equal to one if the child adheres to the vaccination schedule. We construct separate sibling groups based on each parent and assign weight 0.5 for children that appear twice so that each child contributes equally to the estimate regardless of whether one or two parents are observed. Standard errors are clustered by parent.

For non-vaccine adverse events we find no negative effects on uptake among children's vaccines for children born after a parent reports an adverse event. In particular, the estimates for the MMR vaccine and the pneumococcus vaccine are precisely estimated and close to zero. When we restrict the analysis to adverse events from vaccines we find large negative estimates that are statistically significant for the MMR vaccine and close to significant for the vaccine against pneumococcal disease. Although we lack statistical power, these findings, together with the findings for children of individuals with Pandemrix-induced narcolepsy, suggest that experiencing adverse events from vaccines causes individuals to be more reluctant vaccinating their children with well-established and safe vaccines.

Table 3: General adverse events and Children’s Vaccines

	Non-vaccine Adverse Events				Vaccine Adverse Events			
	Measles	Diphtheria	Pneumococcal	Index	Measles	Diphtheria	Pneumococcal	Index
Treated	0.009 (0.016)	0.018 (0.017)	0.006 (0.019)	0.011 (0.015)	-0.102 (0.050)	-0.028 (0.068)	-0.100 (0.063)	-0.077 (0.057)
Dep var. Mean	0.92	0.91	0.85		0.92	0.91	0.85	
Num.Obs.	2,054,933	2,054,933	2,054,933	2,054,933	2,054,933	2,054,933	2,054,933	2,054,933
R2	0.761	0.764	0.744	0.769	0.761	0.764	0.744	0.769
FE: Birth year	X	X	X	X	X	X	X	X
FE: Birth Order	X	X	X	X	X	X	X	X
FE: Parent	X	X	X	X	X	X	X	X

*Notes:* This table displays estimates of  $\beta$  from Equation (6). The dependent variable is a binary indicator equal to one if a child has received the recommended number of vaccine doses. Vaccine adverse events refers to adverse events reported from vaccines (ICD-10-SE J07), excluding adverse events from Pandemrix and COVID-19 vaccines. General adverse events refer to adverse events reported from other types of drugs. Standard errors are clustered by parent and child.

## 10 Fact vs Fallacy—The case of Measles Vaccines and Autism

A remaining question is whether the effects documented above require a scientifically validated link between treatment and adverse event. To answer this question, we isolate a scenario in which an adverse event is likely perceived to result from medical treatment despite lacking scientific recognition, exploiting the MMR vaccine controversy to study younger siblings of children who developed autism soon after receiving the measles vaccine.

In 1998, a study of 12 children was published in *The Lancet*, suggesting a potential link between the MMR vaccine and bowel disease, as well as autism. Based on the article, lead author Andrew Wakefield went on to publicly argue against the continued use of the triple MMR vaccine. Media coverage at the time played a major role in disseminating the suspected link to the public. A 2004 journalistic investigation exposed a serious conflict of interest, as the lead author profited from the demotion of the MMR vaccine (Deer, 2020). The article was retracted in 2010 after the General Medical Council revealed data manipulation and nontransparent case selection. Since then, several large-scale studies have found no causal association between the MMR

vaccine, bowel disease, and autism.<sup>14</sup>

We are interested in the role of personal exposure to autism. Similar to the previous section, we exploit within-family variation by comparing children born before and after an autism diagnosis in the same family. We display results for autism together with results for other mental and behavioral disorders (ICD-10 chapter V) and other common diseases developed before the age of 5 in Figure 10.

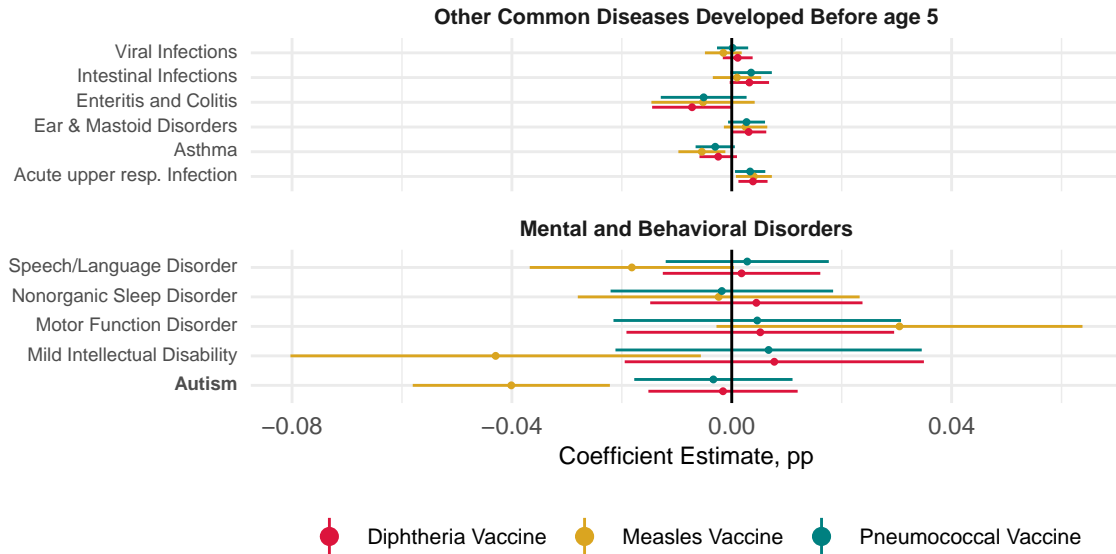


Figure 10: Vaccination Uptake After an Older Sibling Is Diagnosed With Autism

*Notes:* This figure presents estimated coefficients from regressing a binary vaccination adherence indicator on birth order and parent fixed effects, and a binary variable equal to one if the individual was born after an older sibling was diagnosed with autism or another specified disease. The sample consists of individuals born between 2013 and 2022. For autism, the number of treated individuals, defined as those born after an older sibling is diagnosed with autism, is 6,347. Standard errors are clustered by parent.

The findings suggest that children with an older sibling with autism are 4 percentage points less likely to receive the measles vaccine when they are supposed to (a 40% relative increase in non-vaccination compared a baseline 90% vaccination rate), but not other common childhood vaccines. We find a similar pattern for mild intellectual disability, which arguably shares several features with the symptom profile of autism. Once again, the effects are domain specific in the sense that parents do not abstain from the diphtheria vaccine or

<sup>14</sup>It is well established that the Wakefield article lead to increased and lasting skepticism towards the vaccine. For example [Motta and Stecula \(2021\)](#) find a surge in reported adverse events of the vaccine in the US following the publication in 1998.

the pneumococcal vaccine. Note, that we study effects among individuals well after the controversial article linking the MMR-vaccine to autism was retracted in 2010. Our results suggest that even after the retraction of the Wakefield article in 2010, many parents continue to attribute autism to the measles vaccine. This highlights that suspected side effects can persist in public memory irrespective of their scientific foundation: what matters is not scientific consensus, but the presence of a salient cue linking a treatment to a symptom. In line with our broader framework, families exposed to suspected side effects appear to discount official information and instead overweight their own experiences.

## **11 Discussion**

This paper contributes to a growing literature in behavioral economics on experience-based decision making. We study the role of previous experiences in a high-stakes decision where both personal health and others' health are at stake. We provide evidence that adverse drug events affect future healthcare consumption. We use novel and rich Swedish register data on adverse drug reactions and vaccine uptake to measure how narcolepsy, a severe neurological disease linked to the swine flu vaccine in 2010, affected immunization choices during the COVID-19 pandemic. We find that the exposure to narcolepsy lead to significantly lower COVID-19 vaccine uptake, the uptake of fewer doses, and delay in vaccination for the diseased individuals and their close family members. We attribute these findings on vaccine hesitancy to an increase in expected costs of vaccinations through updating of the perceived risks of developing an adverse event. Taken together, the effects of experiencing adverse events are sticky in that they affect the diseased individual a long time after experiencing an adverse event, but are specific in that the experiences only inform a narrow set of similar future decisions that the diseased individual face.

The effects for individuals with high health literacy and for well established vaccines lead us to conclude that negative experiences crowd out official information also when the information is easily accessible. Overall, the results for representative and recurring adverse events suggest that the social cost of adverse events in terms of reduced future uptake of treatment against infectious diseases are limited.

The results on autism and measles vaccinations suggest that public health authorities need to be cautious about the communication of the risk adverse drug events, as false narratives may become part of public memory even if they are not backed-up by scientific evidence. The lack of heterogeneity in our results with respect to health literacy suggests that purely informational nudges about the risk of adverse events may have limited impact in counteracting the negative effects among those that have experienced adverse events.

Against this background, a natural question is what can restore trust and improve immunization outcomes. The role of compensating individuals that have had negative experiences has not been studied. Another interesting venue is the role of positive experiences (as opposed to the valence-neutral healthcare experiences that are considered in this paper) in interfering and undoing negative experiences similar to how some psychologists think about designing cognitive therapy ([Holmes et al., 2006](#); [Ashraf et al., 2024](#)).

## References

- Abadie, Alberto and Jann Spiess (2022) “Robust Post-Matching Inference,” *Journal of the American Statistical Association*, 117 (538), 983–995, [10.1080/01621459.2020.1840383](https://doi.org/10.1080/01621459.2020.1840383).
- Adhikary, Dipannita, Shanto Barman, Redoy Ranjan, and Hana Stone (2022) “A systematic review of major cardiovascular risk factors: a growing global health concern,” *Cureus*, 14 (10).
- Ahmed, S. Sohail, Peter H. Schur, Noni E. MacDonald, and Lawrence Steinman (2014) “Narcolepsy, 2009 A(H1N1) pandemic influenza, and pandemic influenza vaccinations: What is known and unknown about the neurological disorder, the role for autoimmunity, and vaccine adjuvants,” *Journal of Autoimmunity*, 50, 1–11, <https://doi.org/10.1016/j.jaut.2014.01.033>.
- Alsan, Marcella and Marianne Wanamaker (2017) “Tuskegee and the Health of Black Men\*,” *The Quarterly Journal of Economics*, 133 (1), 407–455.
- Anderberg, Dan, Arnaud Chevalier, and Jonathan Wadsworth (2011) “Anatomy of a health scare: education, income and the MMR controversy in the UK,” *Journal of Health Economics*, 30 (3), 515–530.
- Archibong, Belinda and Francis Annan (2023) “‘We Are Not Guinea Pigs’: The Effects of Disclosure of Medical Misconduct on Vaccine Compliance,” NBER Working Paper w31655, National Bureau of Economic Research, [10.3386/w31655](https://doi.org/10.3386/w31655).
- Ashraf, Nava, Gharad Bryan, Alexia Delfino, E Holmes, Leonardo Iacovone, C Meyer, and A Pope (2024) “Learning to see a world of opportunities,” Technical report, Working Paper. Accessed November 22, 2024. [https://navaashraf.com/wp-content/uploads/2024/09/learning\\_to\\_see\\_a\\_world\\_of\\_opportunities\\_1.pdf](https://navaashraf.com/wp-content/uploads/2024/09/learning_to_see_a_world_of_opportunities_1.pdf).
- Barateau, Lucie, Régis Lopez, and Yves Dauvilliers (2016) “Treatment options for narcolepsy,” *CNS Drugs*, 30 (5), 369–379.
- Baumann, Christian R, Emmanuel Mignot, Gert Jan Lammers et al. (2014) “Challenges in diagnosing narcolepsy without cataplexy: a consensus statement,” *Sleep*, 37 (6), 1035–1042.
- Bordalo, Pedro, Giovanni Burro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer (2024) “Imagining the Future: Memory, Simulation, and Beliefs,” *The Review of Economic Studies*, rdae070, [10.1093/restud/rdae070](https://doi.org/10.1093/restud/rdae070).

- Bordalo, Pedro, John J Conlon, Nicola Gennaioli, Spencer Y Kwon, and Andrei Shleifer (2022) “Memory and Probability\*,” *The Quarterly Journal of Economics*, 138 (1), 265–311, [10.1093/qje/qjac031](https://doi.org/10.1093/qje/qjac031).
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2012) “Salience Theory of Choice Under Risk,” *The Quarterly Journal of Economics*, 127 (3), 1243–1285, [10.1093/qje/qjs018](https://doi.org/10.1093/qje/qjs018).
- Bronnenberg, Bart J., Jean-Pierre H. Dubé, and Matthew Gentzkow (2012) “The Evolution of Brand Preferences: Evidence from Consumer Migration,” *American Economic Review*, 102 (6), 2472–2508, [10.1257/aer.102.6.2472](https://doi.org/10.1257/aer.102.6.2472).
- Chang, Lenisa V (2018) “Information, education, and health behaviors: Evidence from the MMR vaccine autism controversy,” *Health economics*, 27 (7), 1043–1062.
- Deer, Brian (2020) *The doctor who fooled the world: Science, deception, and the war on vaccines*: JHU Press.
- Edwards, Kathryn, Germaine Hanquet, Steve Black et al. (2019) “Meeting report narcolepsy and pandemic influenza vaccination: What we know and what we need to know before the next pandemic? A report from the 2nd IABS meeting,” *Biologicals*, 60, 1–7, <https://doi.org/10.1016/j.biologicals.2019.05.005>.
- Folkhälsomyndigheten (2024) “Barnvaccinations programmet i Sverige 2024,” Technical report.
- Gauffin, Helena, Inger Boström, Shala Ghaderi Berntsson, Anna Kristoffersson, Mats Fredrikson, and Anne-Marie Landtblom (2024) “Characterization of the Increase in Narcolepsy following the 2009 H1N1 Pandemic in Sweden,” *Journal of Clinical Medicine*, 13 (3), 652.
- Gauffin, Helena, Thomas Fast, Anastasija Komkova, Shala Berntsson, Inger Boström, and Anne-Marie Landtblom (2022) “Narcolepsy treatment in Sweden: An observational study,” *Acta Neurologica Scandinavica*, 145 (2), 185–192.
- Gilboa, Itzhak and David Schmeidler (1995) “Case-Based Decision Theory,” *The Quarterly Journal of Economics*, 110 (3), 605–639, <http://www.jstor.org/stable/2946694>.
- Giulietti, Corrado, Michael Vlassopoulos, and Yves Zenou (2023) “When reality bites: Local deaths and vaccine take-up,” *European Economic Review*, 156, 104463, <https://doi.org/10.1016/j.eurocorev.2023.104463>.
- Graeber, Thomas, Christopher Roth, and Florian Zimmermann (2024) “Stories, Statistics, and Memory\*,” *The Quarterly Journal of Economics*, 139 (4), 2181–2225, [10.1093/qje/qjae020](https://doi.org/10.1093/qje/qjae020).

- Hallberg, Pär, Hans Smedje, Niclas Eriksson et al. (2019) “Pandemrix-induced narcolepsy is associated with genes related to immunity and neuronal survival,” *EBioMedicine*, 40, 595–604.
- Hartonen, Tuomo, Bradley Jermy, Hanna Sõnajalg et al. (2023) “Nationwide health, socio-economic and genetic predictors of COVID-19 vaccination status in Finland,” *Nature Human Behaviour*, 7 (7), 1069–1083, [10.1038/s41562-023-01591-z](https://doi.org/10.1038/s41562-023-01591-z).
- Holmberg, Sören and Per Hedberg (2020) “Vilka vaccinerade sig mot svininfluensan 2009/10?” Technical report, The SOM Institute.
- Holmes, Emily A., Andrew Mathews, Tim Dalgleish, and Bundy Mackintosh (2006) “Positive Interpretation Training: Effects of Mental Imagery Versus Verbal Training on Positive Mood,” *Behavior Therapy*, 37 (3), 237–247, <https://doi.org/10.1016/j.beth.2006.02.002>.
- Insani, Widya N, Cate Whittlesea, Hassan Alwafi, Kenneth K C Man, Sarah Chapman, and Li Wei (2021) “Prevalence of adverse drug reactions in the primary care setting: A systematic review and meta-analysis,” *PLOS ONE*, 16 (5), e0252161, [10.1371/journal.pone.0252161](https://doi.org/10.1371/journal.pone.0252161), eCollection 2021.
- Komagamine, Junpei (2024) “Prevalence of urgent hospitalizations caused by adverse drug reactions: a cross-sectional study,” *Scientific Reports*, 14 (1), 6058, [10.1038/s41598-024-56855-z](https://doi.org/10.1038/s41598-024-56855-z).
- LaBar, Kevin S. and Roberto Cabeza (2006) “Cognitive neuroscience of emotional memory,” *Nature Reviews Neuroscience*, 7 (1), 54–64, [10.1038/nrn1825](https://doi.org/10.1038/nrn1825).
- Larson, Heidi J., Emmanuela Gakidou, and Christopher J.L. Murray (2022) “The Vaccine-Hesitant Moment,” *New England Journal of Medicine*, 387 (1), 58–65, [10.1056/NEJMr2106441](https://doi.org/10.1056/NEJMr2106441).
- Larson, Heidi J, Caitlin Jarrett, Elisabeth Eckersberger, David MD Smith, and Pauline Paterson (2014) “Understanding vaccine hesitancy around vaccines and vaccination from a global perspective: a systematic review of published literature, 2007–2012,” *Vaccine*, 32 (19), 2150–2159.
- Lowes, Sara and Eduardo Montero (2021) “The Legacy of Colonial Medicine in Central Africa,” *American Economic Review*, 111 (4), 1284–1314.
- Mahlis, Josh, Alberto K De la Herrán-Arita, and Emmanuel Mignot (2013) “The autoimmune basis of narcolepsy,” *Current Opinion in Neurobiology*, 23 (5), 767–773, <https://doi.org/10.1016/j.conb.2013.04.013>, Circadian rhythm and sleep.

- Malmendier, Ulrike (2021a) “Experience effects in finance: Foundations, applications, and future directions,” *Review of Finance*, 25 (5), 1339–1363.
- (2021b) “FBBVA lecture 2020 exposure, experience, and expertise: Why personal histories matter in economics,” *Journal of the European Economic Association*, 19 (6), 2857–2894.
- Malmendier, Ulrike and Stefan Nagel (2011) “Depression babies: do macroeconomic experiences affect risk taking?” *The Quarterly Journal of Economics*, 126 (1), 373–416.
- Malmendier, Ulrike, Stefan Nagel, and Zhen Yan (2021) “The making of hawks and doves,” *Journal of Monetary Economics*, 117, 19–42, <https://doi.org/10.1016/j.jmoneco.2020.04.002>.
- Martinez-Bravo, Monica and Andreas Stegmann (2021) “In Vaccines We Trust? The Effects of the CIA’s Vaccine Ruse on Immunization in Pakistan,” *Journal of the European Economic Association*, 20 (1), 150–186.
- Mereckiene, J, S Cotter, J T Weber et al. (2012) “Influenza A(H1N1)pdm09 vaccination policies and coverage in Europe,” *Euro Surveill.*, 17 (4).
- Motta, Matthew and Dominik Stecula (2021) “Quantifying the effect of Wakefield et al.(1998) on skepticism about MMR vaccine safety in the US,” *PloS one*, 16 (8), e0256395.
- Murphy, Jamie, Frédérique Vallières, Richard P Bentall et al. (2021) “Psychological characteristics associated with COVID-19 vaccine hesitancy and resistance in Ireland and the United Kingdom,” *Nature communications*, 12 (1), 29.
- Ortiz, Luis E., Anne Marie Morse, Lois Krahn, Maggie Lavender, Matthew Horsnell, Dianna Cronin, Beth Schneider, and Jennifer Gudeman (2025) “A Survey of People Living with Narcolepsy in the USA: Path to Diagnosis, Quality of Life, and Treatment Landscape from the Patient’s Perspective,” *CNS Drugs*, 39 (Suppl 1), 23–36, [10.1007/s40263-024-01142-8](https://doi.org/10.1007/s40263-024-01142-8), Open access article; funded by Avadel Pharmaceuticals.
- Oster, Emily (2018) “Does disease cause vaccination? Disease outbreaks and vaccination response,” *Journal of Health Economics*, 57, 90–101, [10.1016/j.jhealeco.2017.10.003](https://doi.org/10.1016/j.jhealeco.2017.10.003).
- Partinen, Markku, Birgitte Rahbek Kornum, Giuseppe Plazzi, Poul Jennum, Ilkka Julkunen, and Outi Vaarala (2014) “Narcolepsy as an autoimmune disease: the role of H1N1 infection and vaccination,” *The Lancet Neurology*, 13 (6), 600–613.

- Persson, I., F. Granath, J. Askling, J. F. Ludvigsson, T. Olsson, and N. Feltelius (2014) “Risks of neurological and immune-related diseases, including narcolepsy, after vaccination with Pandemrix: a population- and registry-based cohort study with over 2 years of follow-up,” *Journal of Internal Medicine*, 275 (2), 172–190, <https://doi.org/10.1111/joim.12150>.
- Pirmohamed, Munir, Sally James, Shaun Meakin et al. (2004) “Adverse drug reactions as cause of admission to hospital: prospective analysis of 18,820 patients,” *BMJ*, 329 (7456), 15–19, [10.1136/bmj.329.7456.15](https://doi.org/10.1136/bmj.329.7456.15).
- Shattock, Andrew J., Helen C. Johnson, So Yoon Sim, Austin Carter, Philipp Lambach, and Raymond C.W. Hutubessy (2024) “Contribution of vaccination to improved survival and health: modelling 50 years of the Expanded Programme on Immunization,” *The Lancet*, 403 (10441), P2307–2316.
- Stowe, Julia, Nick Andrews, Paul Gringras, Timothy Quinnell, Zenobia Zaiwalla, John Shneerson, and Elizabeth Miller (2020) “Reassessment of the risk of narcolepsy in children in England 8 years after receipt of the AS03-adjuvanted H1N1 pandemic vaccine: A case-coverage study,” *PLOS Medicine*, 17 (9), 1–17, [10.1371/journal.pmed.1003225](https://doi.org/10.1371/journal.pmed.1003225).
- WHO (2019) “Ten threats to global health in 2019,” <https://www.who.int/news-room/spotlight/ten-threats-to-global-health-in-2019>, [Online; accessed 2025-05-29].

## Appendix A. Background

### A.1 Narcolepsy

Narcolepsy is a chronic neurological disease that significantly disrupts the ability to control sleep-wake cycles. It is characterized by excessive daytime sleepiness (EDS), which is present in all individuals with the condition. Narcolepsy is classified into two types: type 1 and type 2, which differ in their associated symptoms. In type 1 narcolepsy, patients experience both EDS and cataplexy, a sudden loss of muscle tone triggered by strong emotions. This form is typically associated with an autoimmune process that leads to a deficiency in hypocretin, a neuropeptide critical to the regulation of wakefulness, in the central nervous system. In contrast, patients with type 2 narcolepsy exhibit EDS without cataplexy, and the severity of EDS is generally less pronounced compared to type 1 (see [Baumann et al., 2014](#)).

The clinical diagnosis entails detailed measurement of hypocretin concentration to assess the potential for type 1 narcolepsy. Furthermore, the patient will undergo lengthy examinations in sleep laboratories (multiple sleep latency tests) and several diagnostic tests for exclusion of other drivers of sleep disturbance (ie, analysis of blood plasma for iron deficiency). Due to the observability of cataplexies and established biomarkers available as diagnostic anchors, studies have found that type 1 narcolepsy to be detected more easily by medical practitioners.

Although there is currently no cure to the disease, treatment of symptoms using behavioral measures (i.e., sleep hygiene and schedule) and medical treatment ([Barateau et al., 2016](#)) is available. Patients with narcolepsy often develop conditions, such as metabolic and cardiovascular diseases, psychiatry diseases, musculoskeletal chronic pain, and other specific sleep disorders ([Barateau et al., 2016](#)).

The existing literature suggests that narcolepsy cases are associated with genetic predispositions, such as carrying the genom HLA-DQB1\*06:02 and GDNF-AS1 ([Hallberg et al., 2019](#); [Gauffin et al., 2022](#)). Around 30% of the Swedish and Finnish general population carry the respective alleles and therefore it is not sufficient to explain the development of narcolepsy ([Partinen et al., 2014](#)).

### A.2 Legal process

Patients that developed narcolepsy after vaccination with the Pandemrix vaccine first claimed reimbursement from the Swedish Pharmaceutical Insurance, an insurance that is an agreement between almost all pharmaceutical companies operating in Sweden. The Swedish Pharmaceutical Insurance was obliged to pay out at most 150 million SEK (the common limit to payouts for all injuries attributable within one calendar year). In 2016, the government decided to assume responsibility for compensating affected individuals who had not received adequate compensation from the insurance company, owing to its financial limitations. Individuals received at most 10 million SEK (900k USD) from the insurance company and the government combined—where the purpose of the payment is compensation for losses in salary. The insurance company and the government made different judgements on the requirements for receiving reimbursements but in general individuals aged above 20 at the time of the first vaccination and individuals who reported narcolepsy after three years of vaccination were less likely to receive reimbursement. Ten years after the swine flu vaccine campaign, some diseased individuals are still claiming and receiving reimbursements from the government for pain and suffering.

### A.3 Swine flu pandemic and the vaccination campaign

In April 2009, the first cases of swine flu were discovered in Mexico. In June, three months later, the World Health Organization declared swine flu influenza a pandemic. Phase 1 clinical trials for Pandemrix were completed in September 2009, at which point Pandemrix was granted market authorization by the European Commission given the exceptional circumstances.

During the vaccination campaign, Sweden had the goal of reaching herd immunity. The general public was recommended to take the vaccine and regional healthcare authorities facilitated vaccine campaigns in schools. However, some risk groups were prioritized including individuals with some chronic diseases, pregnant individuals, and healthcare workers, were recommended taking the swine flu vaccine. The vaccine was fully subsidized.

In total 60% of the population received the swine flu vaccine, which was lower than the authorities had initially hoped. Surveys conducted at the time indicates that many were skeptical due to the risk of adverse events already before the vaccination campaign. Another potential reason is that the swine flu spread slower than initially anticipated.

The Swedish Public Health Authority estimates that the vaccination campaign saved around 100 lives and prevented 215 intensive care unit treatments.

The authorization of Pandemrix was fast-tracked due to exceptional circumstances, which allowed it to market with only data from the first phase of clinical trial in place. This was possible because of a mock-up vaccine route, where the complete vaccine protocol, including the adjuvant, had already been tested extensively, which only required the virus strain to be adjusted for development. Starting in May 2009 Glaxosmithkline (GSK) received orders from several countries to supply a vaccine against the swine flu. Sweden had already signed a public procurement contract with GSK in November 2007 in place, stating that they would deliver influenza vaccine in the event of a new pandemic. The contract also prescribed that GSK would not be held responsible for potential adverse events.


## Appendix B. Additional Descriptive Statistics & Results

**Table B1:** Disentangling Selection into Treatment from Selection into Narcolepsy


	Treated	Narcolepsy diagnosis	Pandemrix	Heart attack
Socioeconomic	0.52	0.55	0.53	0.58
Socioeconomic + Health	0.55	0.54	0.53	0.61

*Notes:* This table displays AUC-values for predicted probabilities. Probabilities are computed using XGBoost. First column compares individuals who received Pandemrix and developed narcolepsy to those who received Pandemrix but did not develop narcolepsy. Second column compares individuals who are diagnosed with narcolepsy after 2016 and who hence were unlikely to develop it from Pandemrix to random individuals of the Swedish population. Third column compares individuals who take Pandemrix to individuals residing in the same regions. Fourth column compares individuals who are diagnosed with heart attack to random individuals of the Swedish population. The prediction exercises are run on samples that are balanced in birth year.


### Vaccination Rate



### Number of Doses

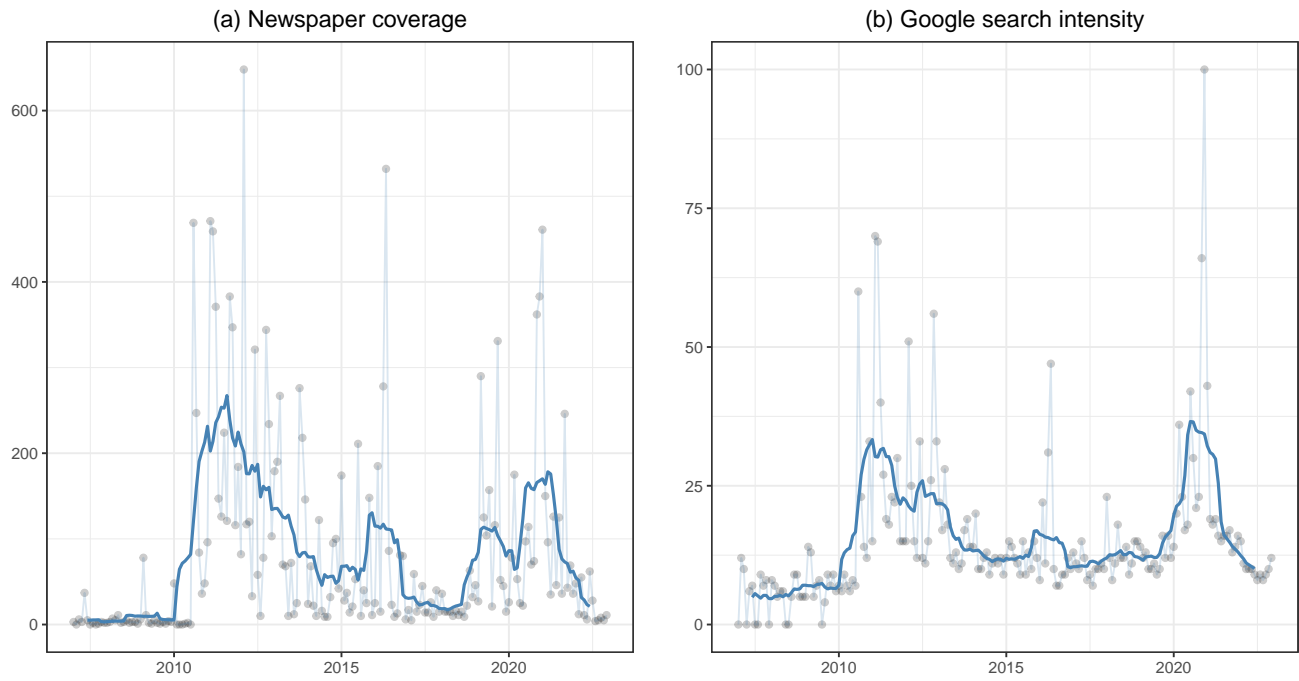


### Days Elapsed



**Figure B1:** COVID-19 immunization outcomes for different sub-populations.

*Notes:* Each row reports COVID-19 immunization outcomes for different sub-populations. To ensure comparability across groups and to accurately measure income and education, we restrict all samples to individuals aged 40–50 in 2022. **First column:** by education level. **Second column:** by gender. **Third column:** by income percentile. **Fourth column:** by geographic origin (Swedish vs. foreign background, where foreign is defined as being born abroad or having two parents born abroad). Outcomes include vaccination rate (top row), average number of doses (middle row), and average days elapsed since first eligibility (bottom row).




**Figure B2:** Exposure and Attention to the Narcolepsy Episode

*Notes:* Panel (a) Monthly number of articles mentioning "Narkolepsi" (eng. narcolepsy) in Swedish press. Panel (b) Monthly Google search intensity for the word "Narkolepsi" (eng. narcolepsy). Each dot correspond to one month while the thick line displays 12-month moving averages.




**Figure B3: Narcolepsy Drug Prescriptions by Year**


*Notes:* This figure shows the share of individuals prescribed each of twelve drugs commonly used to treat narcolepsy symptoms, by year and by sample group. The ATC subgroup N06 corresponds to psychoanalgetics: N06A refers to antidepressants, while N06B refers to psychostimulants (including ADHD medications and nootropics). *Narcolepsy Post Swine flu* refers to individuals diagnosed with narcolepsy after the 2009–2010 swine flu pandemic, while *Narcolepsy Pre Swine flu* refers to those diagnosed earlier.



(a) Post-Swine Flu Cases




(b) Pre-Swine Flu Cases



(c) Pandemrix Vaccinations

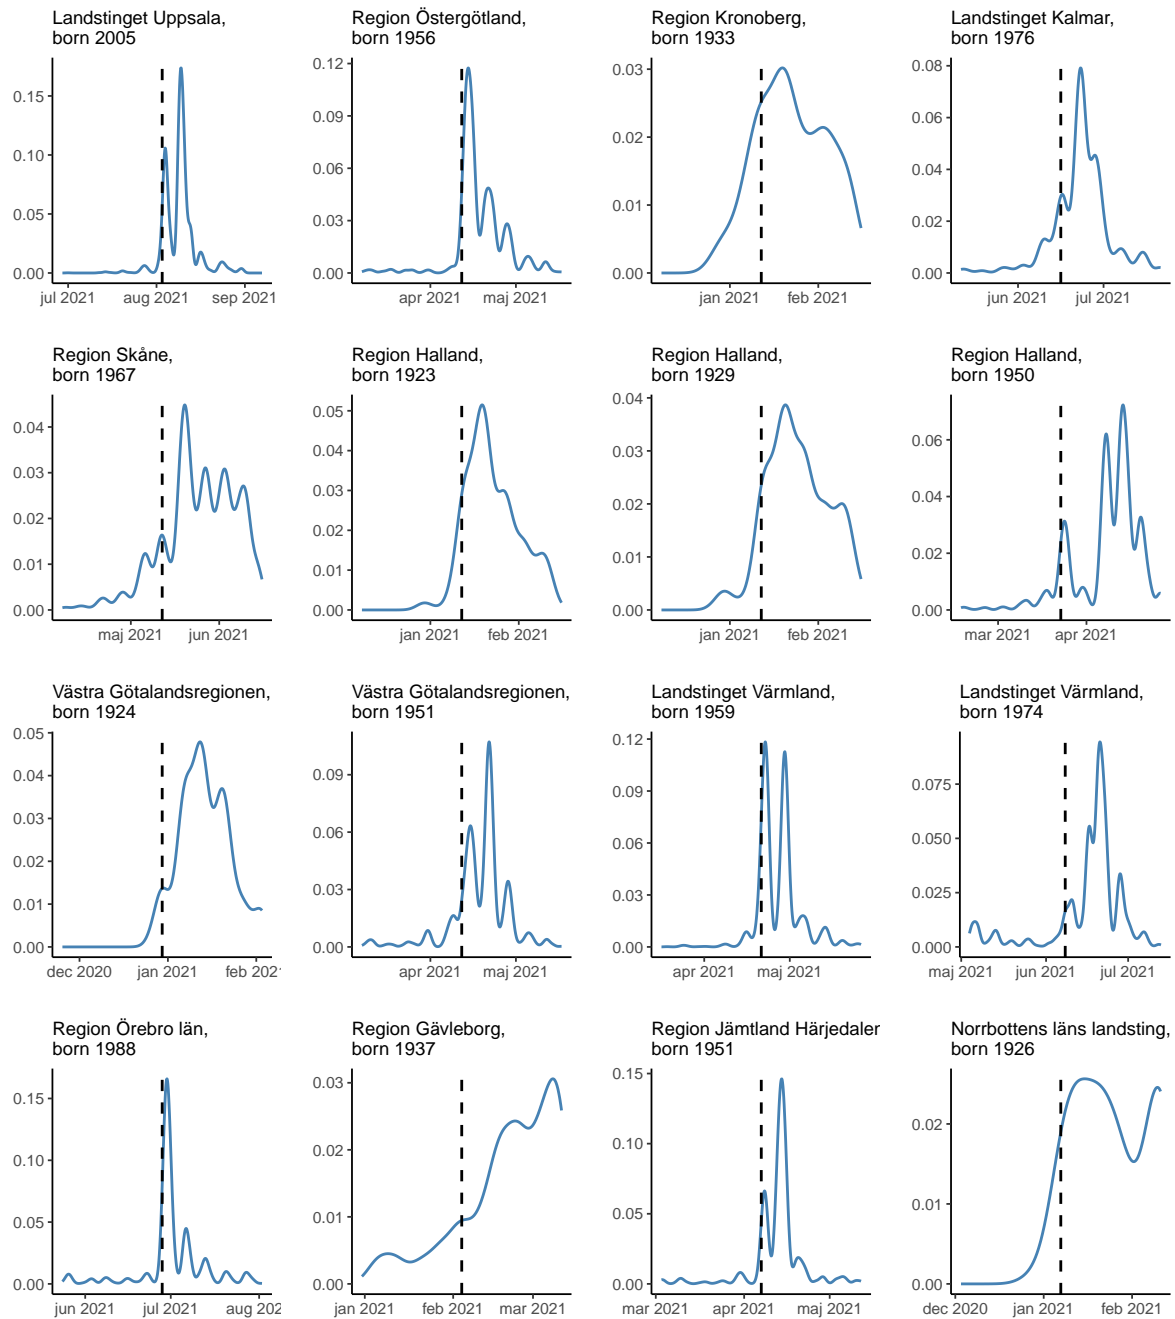
**Figure B4:** Geographical Distribution of Individuals with Narcolepsy and Pandemrix.

*Notes:* **panel (a)** Number of treated individuals by healthcare region of residence in 2011. **Panel (b)** Number of treated individuals by healthcare region of residence in 2011 in the pre-swine flu sample. **Panel (c)** Healthcare regions with individual level Pandemrix data.



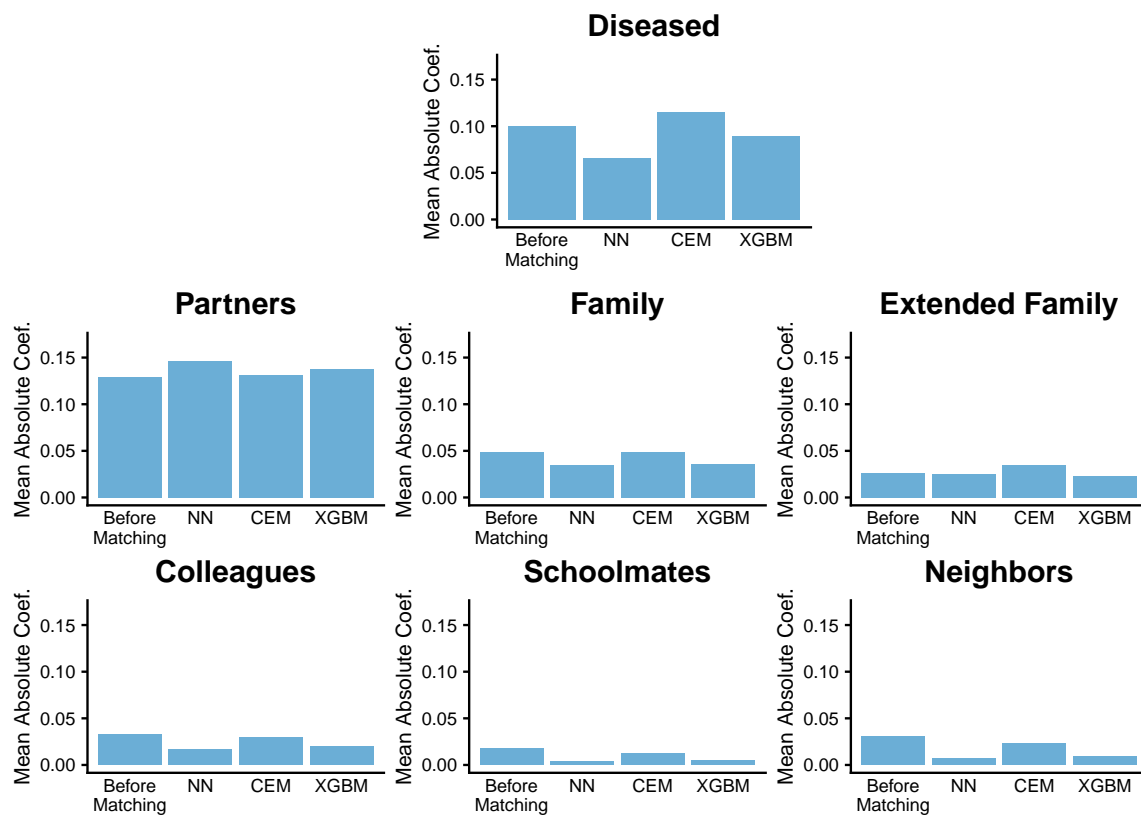
**Figure B5: Balance in Pre-treatment Covariates, Pre-Swine Flu Sample**

*Notes:* This figure displays differences in pre-determined characteristics between individuals that develop narcolepsy and control individuals for the pre sample. Each estimate comes from a univariate regression of the standardized covariate on the treatment status and on birth year fixed effects. Only a subset of covariates are displayed, the full pool of candidate covariates are listed in Section C.




**Figure B8:** Definition of First Date of Availability

*Notes:* This figure displays the distribution of the first date an individual receives the COVID-19 vaccine for a sample of 16 Region x Birth-Year cells, showing a two-month window around the estimated date of vaccine availability (dashed line).



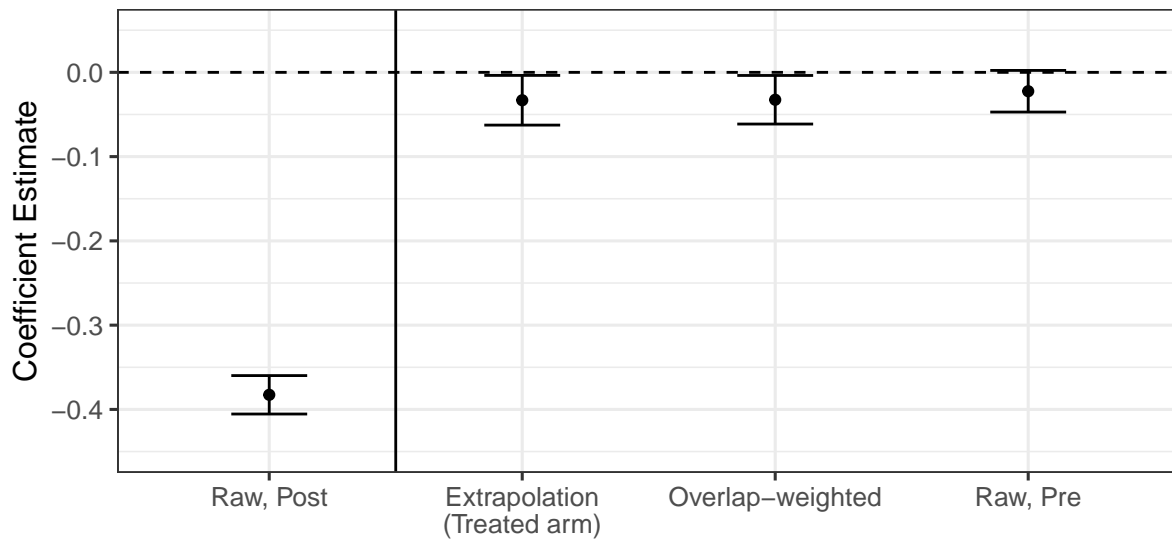
**Figure B6:** Balance Across Matching Methods and Networks, post swine flu sample

*Notes:* This figure shows the mean absolute standardized differences ( $|SMD|$ ) across all covariates between network members to individuals who developed narcolepsy after receiving the swine flu vaccine Pandemrix and network members of individuals who also received Pandemrix but did not develop narcolepsy, for different network definitions and matching methods. **Before matching:** Mean  $|SMD|$  from unadjusted differences, net of birth-year fixed effects. **NN:** Mean  $|SMD|$  after nearest-neighbor matching, where propensity scores are estimated using a logit model with covariate selection via LASSO. **CEM:** Mean  $|SMD|$  after coarsened exact matching on birth year, gender, geographic origin, parental income, parental years of schooling, and parental counts of prescribed drugs and healthcare visits. **XGBM:** Mean  $|SMD|$  after nearest-neighbor matching with propensity scores estimated using gradient boosting.



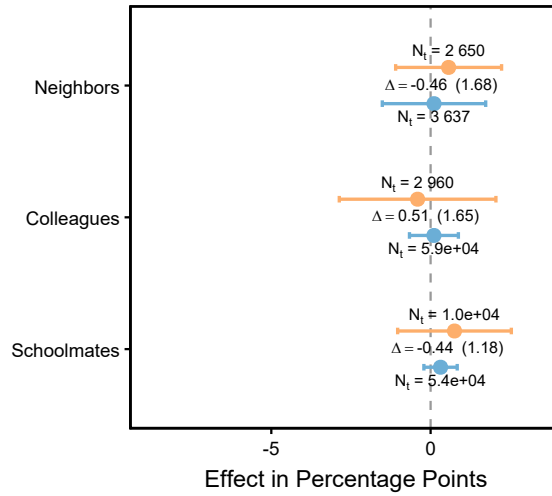
**Figure B7:** Balance Across Matching Methods and Networks, pre swine flu sample

*Notes:* This figure shows the mean absolute standardized differences ( $|SMD|$ ) across all covariates between network members to individuals who developed narcolepsy before the swine flu pandemic and network members of a random sample of 1,000,000 individuals living in Sweden in 2021 but did not develop narcolepsy, for different network definitions and matching methods. **Before matching:** Mean  $|SMD|$  from unadjusted differences, net of birth-year fixed effects. **NN:** Mean  $|SMD|$  after nearest-neighbor matching, where propensity scores are estimated using a logit model with covariate selection via LASSO. **CEM:** Mean  $|SMD|$  after coarsened exact matching on birth year, gender, geographic origin, parental income, parental years of schooling, and parental counts of prescribed drugs and healthcare visits. **XGBM:** Mean  $|SMD|$  after nearest-neighbor matching with propensity scores estimated using gradient boosting.




**Figure B9:** Counterfactual Narcolepsy Effect

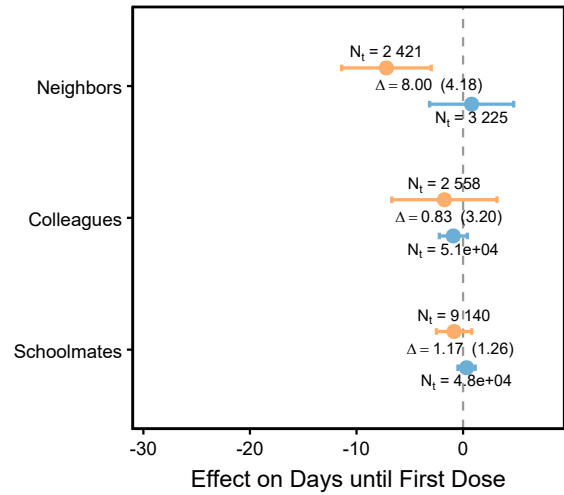
*Notes:* **Raw, Post** is the unadjusted treated–control difference in the main sample. **Extrapolation (treated arm)** is the model–based treatment effect for target ages that lack treated–arm support in the Pre–Swine Flu group. We fit a natural spline in the Pre–Swine Flu group to predict  $Y(1)$  for the age distribution observed among individuals with vaccine–induced narcolepsy. Confidence intervals use 100 bootstrap replications. **Overlap–weighting** is then used to extrapolate the predicted vaccine uptake to the age distribution of post–Swine–Flu individuals based on the Pre–Swine Flu sample. **Raw, Pre** is the unadjusted treated–control difference in the Pre–Swine Flu sample.



(a) Vaccine Taken



(b) Number of Doses

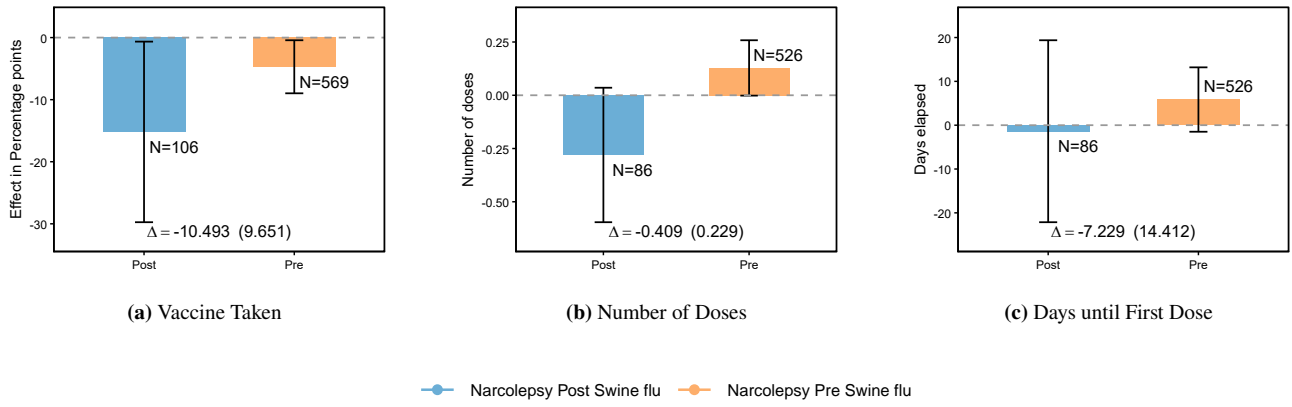


(c) Days until First Dose

● Narcolepsy Post Swine flu ● Narcolepsy Pre Swine flu

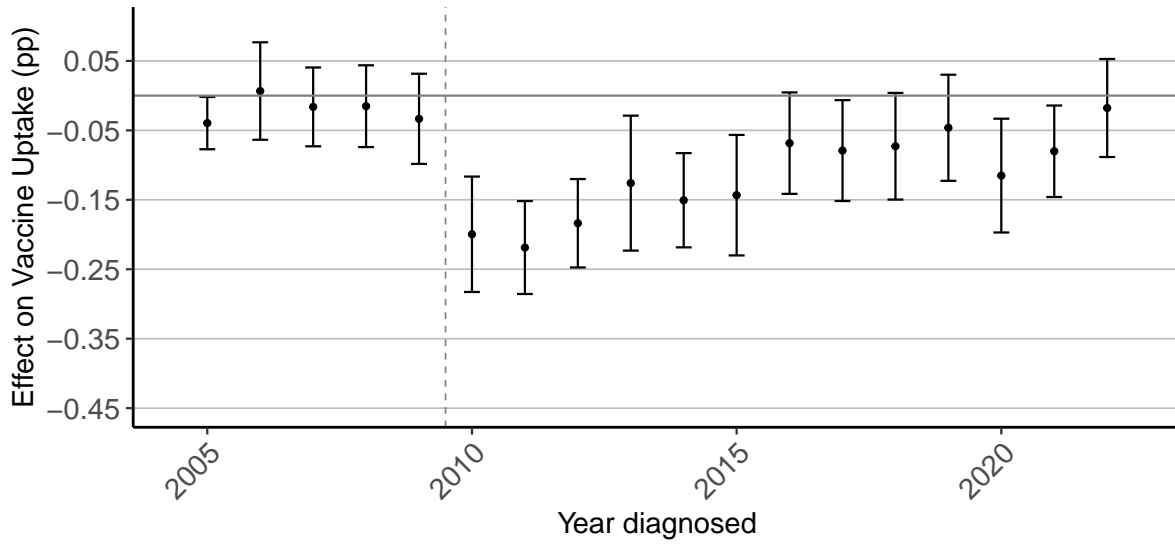
**Figure B10: Main Results—Extended networks.**

*Notes:* This figure displays coefficient corresponding to  $\tau_{post}$  and  $\tau_{pre}$  in eq. (4). The first row corresponds to coefficient for individuals that lived in the same neighborhood as the focal sample individuals in 2011. The second row corresponds to coefficients for colleagues of focal sample individuals in 2021. The third row corresponds to coefficients schoolmates to the focal sample individuals in 2011. **Panel (a):** Displays results for whether or not an individual has taken at least one COVID-19 dose. **Panel (b):** Displays results for the number of COVID-19 doses, conditional on having taken at least one dose. **Panel (c):** Displays results for the number of days elapsed until the first COVID-19 dose is taken. Standard errors are clustered by treatment cluster, defined by the treated or control individual to whom a network member is related as well as at the match level.



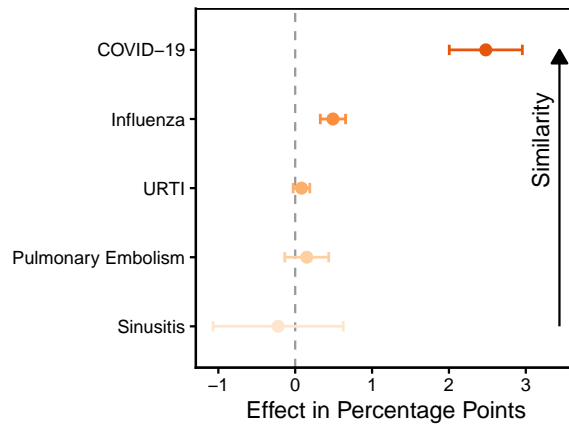
**Figure B11:** Main Results—Partners

*Notes:* This figure displays coefficients corresponding to  $\tau_{post}$  and  $\tau_{pre}$  in eq. (4), similar to fig. 5 but for partners of the focal individuals. **Panel (a):** Whether an individual has taken at least one COVID-19 dose. **Panel (b):** Number of doses, conditional on at least one dose. **Panel (c):** Days elapsed until the first dose. Standard errors are clustered by treatment cluster, i.e. by the treated or control individual that a network member is related to.

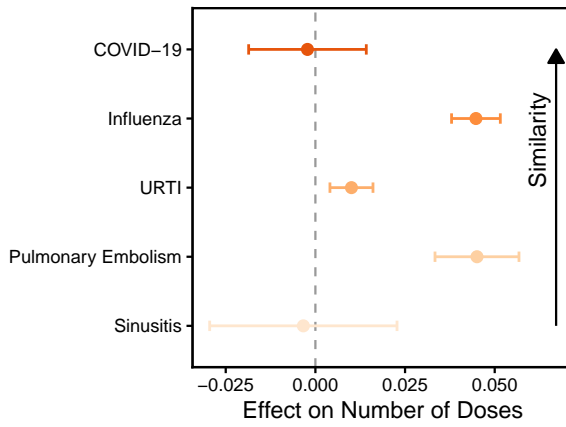


**Figure B12:** Effect on COVID-19 vaccine uptake of Narcolepsy Diagnoses Across Time

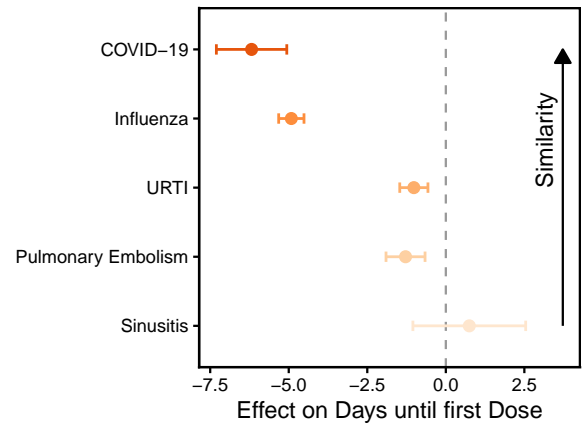
*Notes:* Differences in COVID-19 vaccination rates for different years of first being diagnosed with narcolepsy. The control group for all estimates is a random subset of 1,000,000 individuals living in Sweden. The vertical dashed line marks the onset of Pandemrix vaccinations. All specifications include fixed effects for birth year and gender.



(a) Vaccine Taken



(b) Number of Doses



(c) Days until First Dose

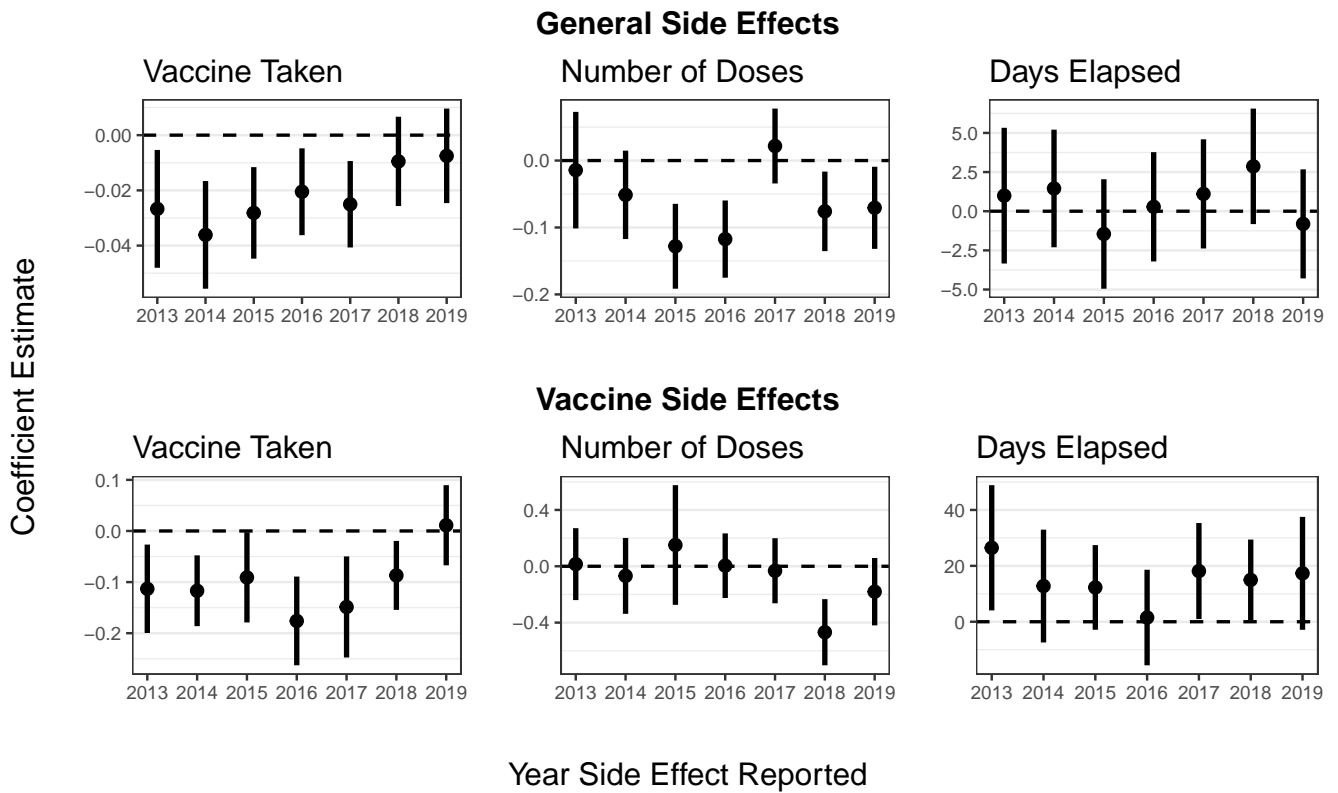
**Figure B13: Benefits of Vaccines—Effects of Prior Diseases on COVID-19 uptake**

*Notes:* **Panel (a):** Displays results for whether or not an individual has taken at least one COVID-19 dose. **Panel (b):** Displays results for the number of COVID-19 doses, conditional on having taken at least one dose. **Panel (c):** Displays results for the number of days elapsed until the first COVID-19 dose is taken. The treatment variable takes the value 1 if an individual was hospitalized with the respective condition listed as the main diagnosis between 2005 and 2021.

**Table B2:** Adverse Events and Immunization Outcomes—Age Reweighted


	Non-vaccine Adverse Events			Vaccine Adverse Events		
	Vaccine Taken	Number of Doses	Days Elapsed	Vaccine Taken	Number of Doses	Days Elapsed
<b>Diseased Individuals</b>						
adverse event	-0.019 (0.003)	-0.025 (0.013)	-1.90 (0.63)	-0.097 (0.013)	-0.062 (0.059)	11.2 (2.95)
N. Treated	15 294	13 974	13 974	849	686	686
<b>Family Members</b>						
adverse event	-0.001 (0.002)	0.048 (0.009)	0.74 (0.38)	-0.042 (0.008)	-0.15 (0.037)	7.83 (1.54)
N. Treated	34 018	31 445	31 445	2 048	1 771	1 771
<b>Children</b>						
adverse event	0.005 (0.003)	0.015 (0.008)	-0.77 (0.52)	-0.029 (0.018)	0.072 (0.044)	3.25 (2.89)
N. Treated	28 087	22 124	22 124	887	662	662
<b>Partner</b>						
adverse event	-0.003 (0.003)	-0.015 (0.018)	-0.78 (0.72)	-0.036 (0.014)	-0.12 (0.086)	5.03 (4.15)
N. Treated	7 430	7 049	7 049	329	311	311

*Notes:* This table displays results from regressing the COVID-19 immunization outcome variables on a variable indicating if someone developed an reported an adverse event between 2015 and 2020. Columns 1–3 use an indicator for reporting any adverse event that was not vaccine-related between 2015 and 2020 as the regressor (“General adverse events”). Columns 4–6 instead use an indicator for reporting a vaccine adverse event (“Vaccine adverse events”). Each treated individual is matched 1:1 to an untreated individual who (a) was born in the same year and (b) took the same drug in the same year as the treated individual. Weights are computed based on birth year using kernel smoothing.



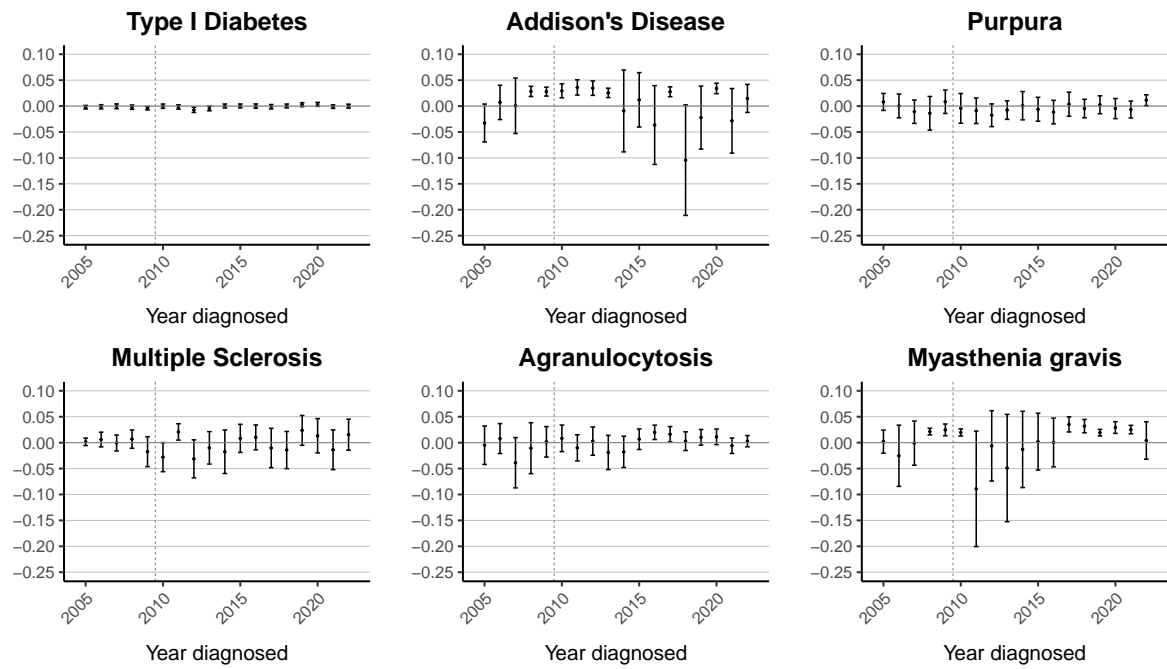
**Figure B14:** Adverse Events and Immunization Outcomes—Heterogeneity by Time.

*Notes:* This figure displays results from regressing the COVID-19 immunization outcomes variables on general adverse events and vaccine adverse events split up by year of reporting the adverse event. Observations are reweighted to match a common age and drug distribution across years.




**Figure B15:** Heterogeneity—Similarity-based learning

*Notes:* This figure displays the estimated coefficients on  $\tau_{post}$  for the outcome variable *Vaccine Taken* among cousins of individuals in the focal sample. First row: Heterogeneity by gender. Second row: Heterogeneity by age, where individuals are considered the same age if their age difference is one year or less. Third row: Heterogeneity based on whether individuals share both gender and age.



**Figure B16:** Other Severe Diseases—Heterogeneity Across Time

*Notes:* This figure shows coefficient estimates from regressions where the outcome variable is Vaccine Taken. Treatment is defined as first being diagnosed with the respective disease in year  $t$ . The control group consists of individuals who received Pandemrix but did not develop narcolepsy. Each specification includes birth-year fixed effects.



**Figure B17:** DiD-Estimates of Healthcare Visits After Narcolepsy Diagnosis

*Notes:* This figure displays difference-in-differences estimates from regressions where the outcome variable is the number of healthcare visits per year. The specification includes a full set of event-time indicators relative to the year of first narcolepsy diagnosis. The sample consists of treated individuals matched 1:1 on birth year to control units who received Pandemrix but did not develop narcolepsy. The light blue line shows coefficients for visits where the main diagnosis is narcolepsy (G47.4, ICD-10-SE), while the dark blue line shows coefficients for visits related to other diagnoses. Standard errors are clustered at the individual level.

**Table B3:** Childhood vaccination—Descriptive Statistics

<b>Panel a: Share adhering</b>				
	Diphtheria	Pneumococcus	Measles	COVID-19 (Parents)
Abstention	1.4	2.3	5.9	6.0
Partial Adherence	4.8	8.7		8.2
Full Adherence	93.9	89.0	94.1	85.8

<b>Panel b: Correlation matrix, Full Adherence</b>				
	Diphtheria	Pneumococcus	Measles	COVID-19 (Parents)
Diphtheria	1.00			
Pneumococcus	0.69	1.00		
Measles	0.34	0.25	1.00	
COVID-19 (Parents)	0.11	0.09	0.13	1.00

*Notes:* The figure displays vaccine uptake within the Swedish childhood vaccination program. The sample includes individuals born between 2013 and 2022, implying that approximately 98% of all vaccine doses will have been administered by the end of the observation period (April 2024). **Panel a:** *Abstention* is defined as having received zero doses of the vaccine. For COVID-19, this corresponds to both parents having received zero doses. *Partial adherence* is defined as having received one or two doses of the diphtheria or pneumococcal vaccine. For COVID-19, this corresponds to at least one parent having received fewer than two doses. At least one parent receiving fewer than two COVID-19 doses. *Full adherence* is defined as having received three or more doses of diphtheria and pneumococcal vaccines and at least one dose of the measles vaccine. For COVID-19, this corresponds to both parents having received two or more doses. **Panel b:** Correlations in full adherence.

**Table B4: Reported Adverse Events**

Non-Vaccine Adverse Events		Vaccine Adverse Events	
Symptom	N. Reports	Symptom	N. Reports
Central nervous system haemorrhages and cerebrovascular accidents	508	Vaccination site reactions	65
Angioedemas	389	Disturbances in consciousness NEC	37
Therapeutic and nontherapeutic responses	385	Paraesthesias and dysaesthesias	30
Nausea and vomiting symptoms	363	Febrile disorders	27
Urticarias	320	Injection site reactions	26
Disturbances in consciousness NEC	306	Musculoskeletal and connective tissue pain and discomfort	25
Anaphylactic and anaphylactoid responses	282	General signs and symptoms NEC	21
Breathing abnormalities	281	Headaches NEC	21
Neurological signs and symptoms NEC	260	Joint related signs and symptoms	21
Rashes, eruptions and exanths NEC	258	Nausea and vomiting symptoms	21
Poisoning and toxicity	235	Neurological signs and symptoms NEC	20
Asthenic conditions	227	Muscle pains	19
Allergic conditions NEC	215	Asthenic conditions	17
Renal failure and impairment	202	Urticarias	16
Hepatobiliary function diagnostic procedures	189	Allergic conditions NEC	15
Others	12 973	Others	440

*Notes:* This table lists the most commonly reported adverse events in the regression sample used in Section 8, excluding those related to influenza vaccines (ATC J07BB02). Column 1 lists non-vaccine-related adverse events, and Column 2 lists vaccine-related events. Repeated reports of the same symptom by the same individual are excluded.

## Appendix C. Covariate Definitions and description of matching methods used

We here define the set of candidate covariates used. Parental variables are defined as either mean, mode or max of the variables displayed in Table C1

**Table C1:** Description of Variables

<b>Variable</b>	<b>Description</b>
<b>Birth year</b>	Year of birth
<b>Days sick</b>	Yearly average number of days reported sick during 2005–2009
<b>Days unemployed</b>	Yearly average number of days unemployed (As defined by swedish Public Employment Service) 2005–2009
<b>Gender</b>	Gender of the individual
<b>Income</b>	Yearly average inflation adjusted disposable income during 2005–2009
<b>Siblings</b>	Number of full, biological siblings
<b>Population density</b>	Population density in individual’s neighborhood in 2009
<b>Geographic Origin</b>	Indicator for either Sweden, Rest of Europe, or Rest of World depending on birth country of the individual and her parents.
<b>Number of drugs</b>	Yearly average number of drugs prescribed during 2005–2009
<b>Number of visits</b>	Yearly average number of specialist healthcare visits during 2005–2009
<b>Years of schooling</b>	Highest level of schooling attained as of 2009 based on a mapping from the SUN classification (Svensk utbildningsnomenklatur) to number of years
<b>Field of Education</b>	Last field of education as of 2009 based on the SUN classification (Svensk utbildningsnomenklatur). 25 broad educational categories, defined based on degrees from high school, upper secondary school or university.
<b>Drugs taken</b>	14 binary variables. One for each level 1 ATC code. Equal to 1 if an individual received a drug within that ATC category at least once between 2005 and 2009.
<b>Diagnoses</b>	21 binary variables. One for each ICD-10 chapter. Equal to 1 if an individual was diagnosed within that Chapter category at least once between 2005 and 2009.

**Nearest Neighbor, Logit (NN)** For the analyses shown in Figure 5, Figure 6, Figure B11, and Figure B10, our preferred approach is 1:1 nearest-neighbor matching with exact matching on gender and birth year. Propensity scores are estimated using a logit model. To allow for different information sets across cohorts, we compute propensity scores separately for individuals born before and after 1990, using distinct pools of candidate matching variables. For those born before 1990, we use 136 candidate variables; for those born after 1990, we exclude own (rather than parental) socioeconomic characteristics, leaving 103 candidates. Likewise, for individuals who developed narcolepsy prior to the swine flu pandemic, we remove variables related to both own socioeconomic and health characteristics, regardless of birth year. Variable selection for the propensity score model follows a LASSO procedure with cross-validation to choose the penalty parameter that minimizes out-of-sample prediction error. Finally, the covariates selected by the LASSO are included in the regression estimated on the matched sample to address any remaining imbalance.

**Nearest Neighbor, Gradient Boosting (XGBM)** We also estimate propensity scores using an XGBoost classifier with a logistic objective function. The model is trained on one-hot–encoded covariates, allowing categorical variables and missing values to be handled internally. Five-fold stratified cross-validation is used, and hyperparameters are kept at standard values for moderate regularization ( $\eta = 0.1$ , max depth = 4, min child weight = 1, subsample = 0.8, fraction of predictors sampled for each tree = 0.8).

**Coarsened Exact Matching (CEM)** We manually choose covariates capturing socioeconomic, health, and demographic characteristics: parental income and years of schooling; parents’ number of prescribed drugs and healthcare visits; and the individual’s gender, age, and region of origin. Because the procedure requires exact matches on these variables, including too many covariates makes it difficult to find suitable matches. Adding additional covariates do not meaningfully improve balance. Bins for numeric variables are created using Sturges’ rule.

## **Appendix D. Regional Pandemrix data**

Table [D1](#) contains information about the reported mild symptoms from Pandemrix. 63% of reported symptoms from Pandemrix are classified as mild, primarily consisting of symptoms like febrile disorders, connective tissue pain, headache, or others. 54% of adverse events are self reported, the remaining reports are made by doctors and nurses. In our preferred specification we include both individuals reported adverse events themselves and individuals that had a doctor or a nurse report the adverse event for them.

**Table D1:** Reported Mild Adverse Events from Pandemrix.

Symptom	Share of reports	Self reported	Professionally reported
General signs and symptoms NEC	0.10	0.70	0.30
Febrile disorders	0.06	0.53	0.47
Musculoskeletal and connective tissue pain and discomfort	0.06	0.99	0.01
Injection site reactions	0.06	1.00	0.00
Pain and discomfort NEC	0.05	0.44	0.56
Headaches NEC	0.05	0.66	0.34
Asthenic conditions	0.04	0.71	0.29
Nausea and vomiting symptoms	0.04	0.53	0.47
Neurological signs and symptoms NEC	0.04	0.60	0.40
Paraesthesias and dysaesthesias	0.04	0.55	0.45
Urticarias	0.03	0.16	0.84
Muscle pains	0.02	0.54	0.46
Joint related signs and symptoms	0.02	0.68	0.32
Feelings and sensations NEC	0.02	0.75	0.25
Rashes, eruptions and exanthems NEC	0.02	0.26	0.74
Others	0.33	0.42	0.58

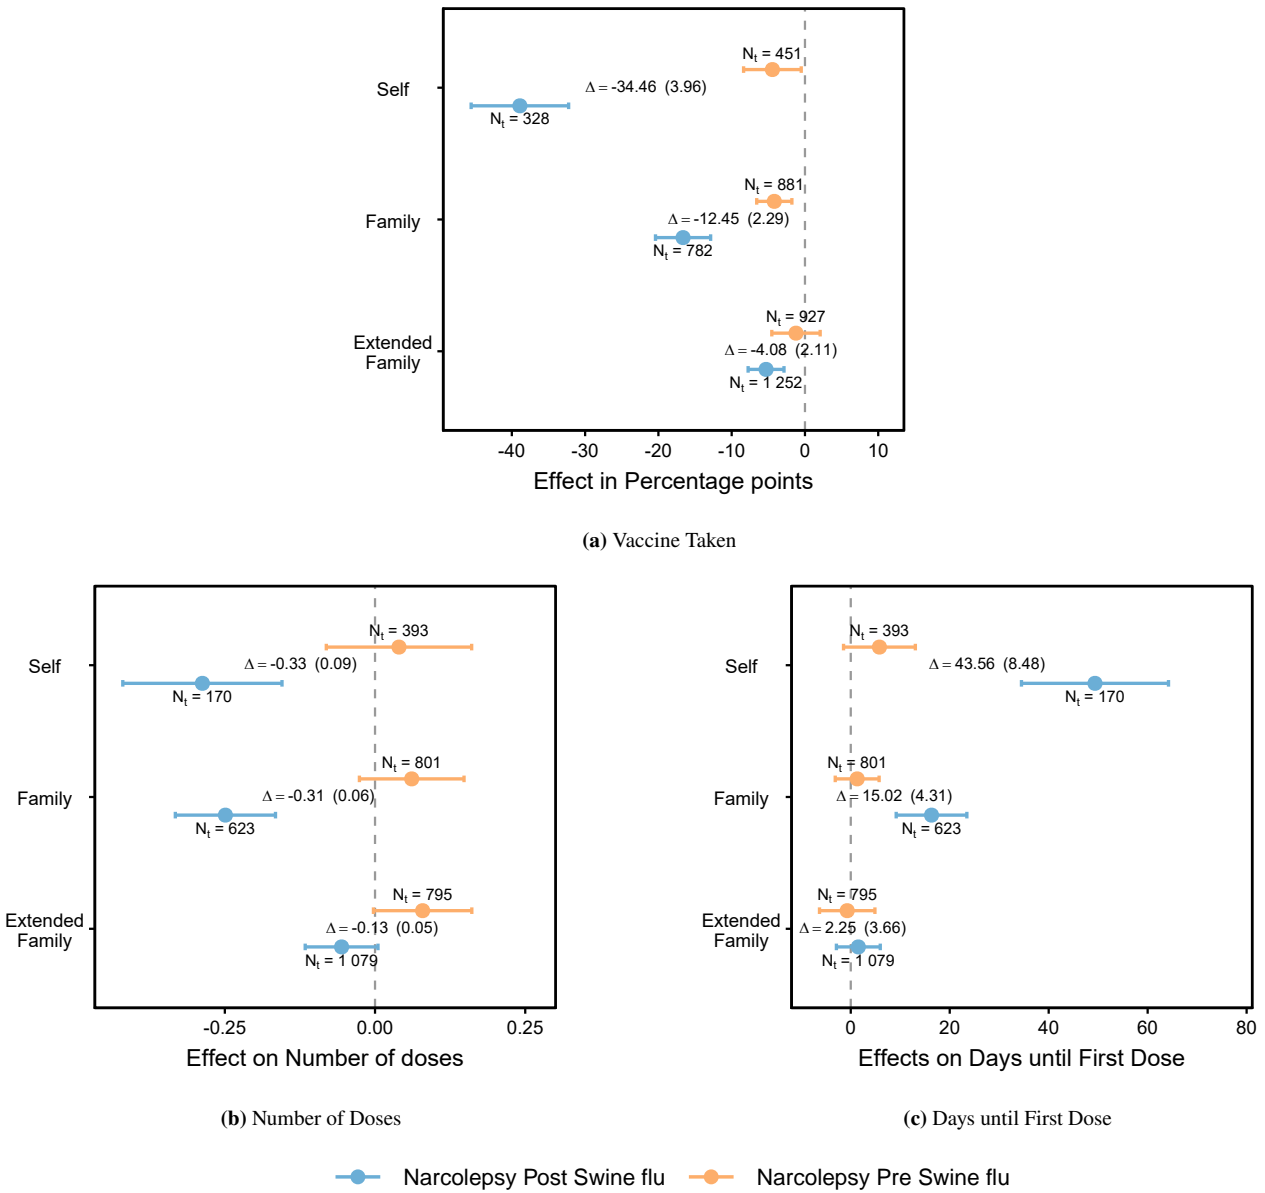
*Notes:* This table displays the most common reported non-severe adverse events by the 2796 individuals who only reported mild symptoms along with the share of each symptom that was self reported and reported by professionals for the same individuals. We remove instances where an individual report the same symptom multiple times.

**Table D2:** Pandemrix Data From Regional Healthcare Authorities

Region	N obs	Vaccination Rate	Comment
Dalarna	167,855	0.61	
Gävleborg	175239	0.63	
Jönköping	231,374	0.69	
Kalmar	78,113	0.33	Covers vaccinations from primary care records only.
Kronoberg	66,488	0.36	Covers vaccinations from primary care records only.
Norrbottn	131,195	0.53	
Uppsala	181,461	0.54	
Värmland	116,931	0.43	
Västerbotten	3,525	0.01	Covers vaccinations after 2010-10-01 only.
Östergötland	274,405	0.64	


*Notes:* This table displays number of individuals receiving Pandemrix by healthcare region from the regions for which we access data. Vaccination rates are defined as the share vaccinated among individuals residing in each healthcare region in 2009.

## Appendix E. Robustness of Main results

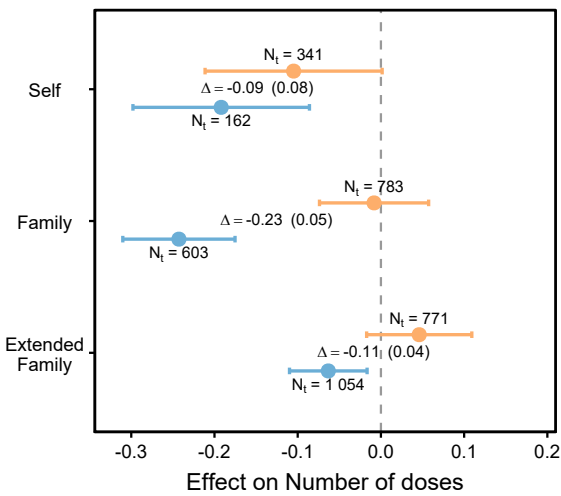


**Figure E1:** Main Results—Diseased Individuals and Family Members, XGBM Propensity Scores

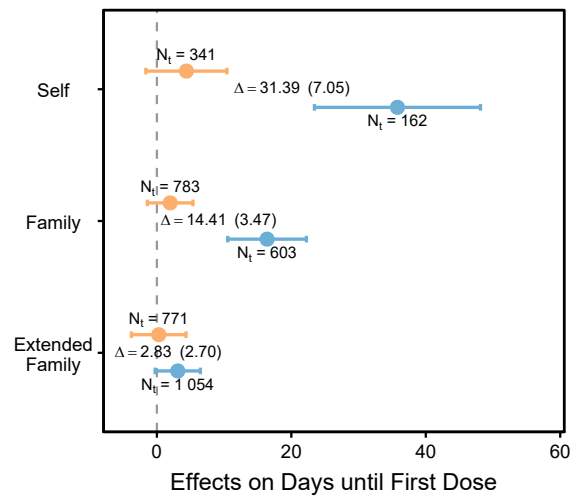
*Notes:* This figure displays coefficients corresponding to  $\tau_{post}$  and  $\tau_{pre}$  in eq. (4). First row corresponds to individuals that themselves developed narcolepsy (Self). The second and third rows show the corresponding coefficients for Family members, defined as siblings and parents, and Extended family members, defined as aunts/uncles and cousins. **Panel (a):** Ever taken at least one COVID-19 dose. **Panel (b):** Number of doses, conditional on at least one dose. **Panel (c):** Days elapsed until the first COVID-19 dose. Standard errors are clustered by treatment cluster, defined by the treated or control individual to whom a network member is related as well as at the match level.



(a) Vaccine Taken



(b) Number of Doses



(c) Days until First Dose

● Narcolepsy Post Swine flu    ● Narcolepsy Pre Swine flu

**Figure E2: Main Results—Diseased Individuals and Family Members, Coarsened Exact Matching**

*Notes:* This figure displays coefficients corresponding to  $\tau_{post}$  and  $\tau_{pre}$  in eq. (4). First row corresponds to individuals that themselves developed narcolepsy (Self). The second and third rows show the corresponding coefficients for Family members, defined as siblings and parents, and Extended family members, defined as aunts/uncles and cousins. **Panel (a):** Ever taken at least one COVID-19 dose. **Panel (b):** Number of doses, conditional on at least one dose. **Panel (c):** Days elapsed until the first COVID-19 dose. Standard errors are clustered by treatment cluster, defined by the treated or control individual to whom a network member is related as well as at the match level.

## Appendix F. Description of matching and predictions

### F.1 Predictions in Table B1

For each prediction, each treated unit is matched to 100 control units with the same birth year. Control units consist of the Swedish population as of 2021.

**Table F1:** XGBoost hyper-parameters, cross-validation setup, and features

<b>Model hyper-parameters</b>	
Maximum tree depth	4
Learning rate ( $\eta$ )	0.10
Minimum child weight	1
Subsample fraction used for training	0.8
Number of boosting rounds	100
<b>Model setup</b>	
Cross-validation folds	5
Number of treated units	1 013
<b>Features</b>	
<b>Socioeconomic:</b> income, years of schooling, field of education, days unemployed, parental income, days on sick leave, parental years of schooling, number of siblings, parental number of siblings, parental days on sick leave, gender, world region of origin.	
<b>Health:</b> 4-digit level drug codes (ATC), 3-digit level diagnosis codes (ICD 10-SE).	

## F.2 Predicting who takes the COVID-19 vaccine

We predict COVID-19 vaccination using a hybrid model that combines a sequence encoder for diagnosis histories with gradient-boosted trees on time-invariant covariates. Each individual’s diagnosis history is converted into an integer sequence, embedded and passed through a bidirectional LSTM, and projected into a low-dimensional representation. The encoder is trained with binary cross-entropy loss and a sigmoid output, using a 10% hold-out split for validation. We concatenate the learned sequence embedding with socioeconomic covariates and use the resulting feature vector to train an XGBoost classifier with logistic loss. Final performance is evaluated on the 10% hold-out test set.

**Table F2:** RNN + XGBoost hyper-parameters, data splits, and features

<b>Sequence encoder (RNN)</b>	
Input	Label-encoded diagnosis codes (with mask_zero)
Vocabulary size	998
Sequence length	402 (truncated at 99th percentile)
Embedding layer	Dimension = 32
Recurrent layer	Bidirectional LSTM, 64 units per direction
Dense projection	32 units, ReLU
Output head	1 unit, sigmoid
Loss / optimizer	Binary cross-entropy / Adam
Batch size / epochs	256 / 5
Validation split	10% hold-out (stratified)
<b>Gradient boosting classifier (XGBoost)</b>	
Objective	binary logistic
Boosting rounds	4000 (with early stopping)
Learning rate ( $\eta$ )	0.02
Max depth	5
Subsample / fraction of predictors sampled for each tree	0.8 / 0.8
Regularization	$\lambda = 1.0, \alpha = 0.0$
Class imbalance	$= n_{\text{neg}} / n_{\text{pos}}$
Early stopping	200 rounds on a 10% validation slice of training data
<b>Model setup</b>	
Train / test split	90% / 10% (stratified)
Hybrid features	Concatenate sequence embedding (32-D) with tabular covariates
<b>Feature sets</b>	
Socioeconomic	income, years of schooling, field of education, days unemployed, parental income, days on sick leave, parental years of schooling, number of siblings, parental number of siblings, parental days on sick leave, gender, world region of origin
Health	4-digit ATC drug codes (sequence input), 3-digit ICD-10-SE diagnoses (sequence input), reported adverse events (sequence input)

## Appendix G. Mathematical Derivations

Observe that MAR among true developers

$$D \perp Y(1) \mid W = 1, X.$$

implies that

$$E[Y(1) \mid W = 1, D = d, X] = E[Y(1) \mid W = 1, X] \quad (d = 0, 1).$$

similarly, MAR among non-developers implies that

$$E[Y(0) \mid W = 0, D = d, X] = E[Y(0) \mid W = 0, X] \quad (d = 0, 1).$$

MAR among developers and source irrelevance (together with the law of total probability) directly give us that

$$\begin{aligned} E[Y \mid D = 1, X] &= E[ WY(1) + (1 - W)Y(0) \mid D = 1, X] \\ &= P(W = 1 \mid D = 1, X) E[Y(1) \mid W = 1, D = 1, X] \\ &\quad + P(W = 0 \mid D = 1, X) E[Y(0) \mid W = 0, D = 1, X] \\ &= E[Y(1) \mid W = 1, D = 1, X] \quad (\text{source irrelevance}) \\ &= E[Y(1) \mid W = 1, X] \quad (\text{MAR, developers}). \end{aligned}$$

What about  $E[Y(0) \mid D = 0, X]$ ? We can write:

$$\begin{aligned} E[Y \mid D = 0, X] &= \Pr(W = 1 \mid D = 0, X) E[Y(1) \mid W = 1, D = 0, X] \\ &\quad + \Pr(W = 0 \mid D = 0, X) E[Y(0) \mid W = 0, D = 0, X]. \end{aligned}$$

That is, it is a weighted average of false negatives and true negatives. Bayes formula yields that.

$$\Pr(W = 1 \mid D = 0, X) = \frac{\Pr(D = 0 \mid W = 1, X) \Pr(W = 1 \mid X)}{\Pr(D = 0 \mid X)}.$$

We maintain the assumption that  $\Pr(W = 1 \mid X)$  is small relative to the probability mass of non-reporters such that the term vanishes. Together with MAR among non-developers it directly follows that

$$E[Y \mid D = 0, X] \approx E[Y \mid D = 0, W = 0, X] = E[Y \mid W = 0, X]$$

## Appendix H. Classifying Adverse Events as Type-Learning or Idiosyncratic

We classify adverse events as either Type-learning or idiosyncratic based on High Level Terms from the MedDRA coding system. These contain very brief descriptions of adverse events, in combination with four character ATC drug codes. We observe 1713 such combinations. To classify these combinations, we use OpenAI’s GPT-4o model with the following system prompt:

I will give you drug side effect combinations and you will classify them as either "idiosyncratic" or "learn about predisposition". Idiosyncratic should be 1 if getting the side effect from that drug does not teach me that I have a predisposition to side effects from COVID-19 mRNA vaccines. The opposite is that experiencing an adverse event teaches me about my likelihood of experiencing an adverse events. Be fairly generous in defining them as idiosyncratic. Output your answer as a JSON array of objects, where each object contains three fields: "drug": ATC code of the drug, "HLT": HLT symptom, "idiosyncratic": 1 or 0, "Justification": Your short justification."

For each combination, we compute an idiosyncrasy score from ten independent model queries using the same input prompt. Each combination is classified as idiosyncratic if its score is above the sample median and as type-learning otherwise. In Table H1, we show three common combinations classified as idiosyncratic and three classified as type-learning.

**Table H1:** Examples of idiosyncratic and type-learning adverse events

<b>Idiosyncratic</b>	<b>Type-learning</b>
<b>Psychostimulants</b> — Fluctuating mood symptoms	<b>Penicillins</b> — Allergic conditions NEC
<b>Aspirin</b> — Nasal disorders (epistaxis)	<b>Enalapril</b> — Angioedemas
<b>Levothyroxine</b> — Asthenic conditions	<b>Sirolimus</b> — Breathing abnormalities

*Notes:* Adverse events are classified as *type-learning* or *idiosyncratic* using MedDRA High Level Terms paired with 4-character ATC drug codes. We define 1,713 HLT–ATC combinations and compute an idiosyncrasy score per combination from ten independent model queries. This table shows three representative examples from each class.