Adverse Health Events and Vaccine Hesitancy

Svenja Miltner* Jonatan Riberth†

[Job Market Paper]

CLICK HERE FOR MOST RECENT VERSION

Abstract

This paper studies how previous experiences with adverse health events affect vaccine hesitancy. Using nationwide Swedish administrative data linking reported adverse drug reactions to vaccination records, we first examine how individuals draw on their own past experiences when making new vaccination decisions. We analyze a severe, well-identified case: narcolepsy, a serious 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; the effects do not attenuate among those with high health literacy or extensive prior healthcare contact. Second, we assess broader social costs of routine adverse events by studying serious events from all licensed pharmaceuticals. Here, people's learning is narrow: adverse events from non-vaccine drugs have little effect on later vaccination, whereas serious vaccine-related events substantially reduce uptake. Together, the two analyses show that individuals rely heavily on their own experiences and that these effects are long lasting but domain-specific. Policy implications follow: while the overall impact of routine adverse events on vaccination is limited, rare but severe vaccine-related events can meaningfully lower uptake.

Keywords: Vaccine hesitancy, Adverse events, Selective memory

JEL Codes: D01, I12, I14

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

†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, Olof Östergren, and seminar participants at IIES, Stockholm University, Freie Universität Berlin, Paris School of Economics, SWECOV Workshop, SUDSWEC Workshop, Université Libre de Bruxelles, SHEA/HEFUU Workshop. 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 Europe has experienced an alarming resurgence in recurring outbreaks of formerly controlled communicable diseases in the past decades. The inability to contain vaccine-preventable diseases is driven mainly 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).

One potential, and arguably the most natural, source of vaccine hesitancy that prior research has largely overlooked is the perceived risk of adverse events. Understanding this channel is important because adverse drug reactions are common: Meta-analyses suggest that 5–10% of hospitalizations are due to adverse drug reactions (Insani et al., 2021; Komagamine, 2024). Beyond their immediate costs in the form of medical treatment and illness, adverse events may result in lower future uptake of vaccines and other medical treatments, imposing additional social costs—in particular through reduced uptake of vaccines against infectious diseases. Assessing the scope of these indirect social costs is essential for policies aiming to increase vaccination uptake and maintain population immunity.

Using administrative records, this paper provides evidence on the effects of experiences with adverse events on future immunization outcomes. While an existing literature links vaccine hesitancy to information gaps, trust in institutions, and peer influence, little is known about how personally experienced adverse drug events shape subsequent demand for health care later in life. Our study is the first to quantify these long-run behavioral responses using data that link 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 plausibly linked to a vaccine deployed during the 2009–2010 swine flu pandemic. During the swine flu pandemic, 60 % of the Swedish population was immunized with the vaccine *Pandemrix* in a national campaign. Epidemiologists later observed increased narcolepsy incidence, a chronic neurological condition with excessive daytime sleepiness and sudden muscle weakness. 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.

In a second part of the analysis we turn to reported severe adverse events from all licensed pharmaceuticals as reported through spontaneous reporting systems[1] used for surveillance purposes. These reports encompass both novel and well-established treatments, making them plausibly representative of the full spectrum of routinely occurring adverse events. This exercise therefore serves to assess the overall impact on subsequent vaccine uptake from adverse events using that are broadly representative of those that patients encounter in real-world practice.

For the first part, we compare COVID-19 vaccination outcomes of individuals who developed narcolepsy after receiving the 2009–2010 swine-flu vaccine with those who received the same vaccine but did not develop narcolepsy, as well as with people who had narcolepsy for other reasons before the campaign. This difference-in-differences design isolates the effect of experiencing a severe, vaccine-induced adverse event on later vaccine hesitancy.[2] For the second part, we rely on adverse events that are reported by healthcare professionals. Furthermore, the data include detailed demographic, socioeconomic, and health information for the entire Swedish population, allowing us to identify comparable individuals not only by demographics and economic background but also by prior health status as reflected in medication use and medical diagnoses. In particular, we compare vaccination outcomes of those experiencing 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 that experiencing a severe adverse events scars individuals with respect to future vaccination

[1] Most high-income countries operate such reporting systems to detect new safety signals and to gauge the frequency of adverse events after a product is introduced.

[2] 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.

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 propensity to vaccinate during the COVID-19 pandemic more than ten years later compared to similar individuals. These effects spill over into personal networks of the affected. Close family members of vaccine-induced narcolepsy patients have a 10 percentage point lower propensity to vaccinate. These findings suggest that individuals are not perfectly informed but instead rely heavily on their own experiences—as opposed to officially provided information—to form an assessment about the *common* risk of experiencing adverse events from the COVID-19 vaccine. By eliciting perceived risks associated with developing COVID-19, or equivalently the benefits of preventive COVID-19 vaccination, we rule out the possibility that the results are driven by changes in the perceived benefits of vaccines rather than by heightened concerns about adverse events. Partners of individuals with narcolepsy show similar patterns to other members family, suggesting the results are not driven by beliefs about genetic predisposition. In addition, individuals who have benefited from vaccinations through exposure to communicable but vaccine-preventable diseases, such as influenza, show positive effects on vaccination outcomes, highlighting that positive experiences also influence uptake. 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. Furthermore, 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 information, even when the information is reliable and accessible.

When we turn to the analysis of recurring and representative adverse events, we find large effects for adverse events from vaccines but small effects for non-vaccine drugs; Experiencing a non-vaccine related adverse event in the ten years preceding the COVID-19 pandemic is associated with a 0.6 percentage-point reduction in COVID-19 vaccine uptake, whereas vaccine-related adverse events cause a 7.5 percentage-point decrease. 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. While severe adverse events are common and their effects persist, generalization is narrow; only similar experiences have a meaningful impact on subsequent decisions. These two observations boil down to the conclusion that the overall impact of adverse events on vaccine hesitancy is limited. However, the results

for vaccine-related adverse events may be 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 usual framework—where individuals rely on personal experiences when domain-specific information is scarce or hard to access—cannot fully explain the reduced uptake. Instead, we interpret these effects on immunization outcomes as suggestive evidence of a shift in the degree to which individuals rely on their own experience as opposed to external information sources; information is easily accessible but individuals do not trust the information. 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-lasting literature on decision making under risk, and in particular, a growing literature on the lasting impact of previous experiences for decision making (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 the selective-memory literature, 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 ensuing scandal. 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 three 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 sporadically 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 more 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) could be paid out for all injuries attributable within one calendar year. 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 assuned the responsibility, (3) 2020 when COVID-19 cases were on the rise. Newspapers published the stories of those affected of the narcolepsy scandal, documented their endeavors to claim reimbursement, and referred back to the scandal when COVID-19 appeared on the public health agenda. The general public partook in the discussion and inquired about the disease online. These observations are meaningful to our analysis in two ways. First, there was at least minimal exposure and therefore knowledge about the narcolepsy scandal 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 scandal, which we implement into 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 five diagnosed cases. 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. 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. 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 this does not make them categorically 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 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 feature high-stakes and inevitably emotionally charged decisions in which individuals are confronted with information from a range of external sources, including authorities, personal networks, and news outlets. With technical and potentially conflicting guidance obscuring clarity, navigating healthcare can be difficult for the layperson.

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. In consequence, when first deciding whether to take the COVID-19 vaccine, people were faced with limited external data on the actual risk of infection and potential vaccine adverse events.

A large literature in economics examines how past experiences shape choices across settings—e.g., high-inflation episodes and stock-market shocks (Malmendier and Nagel, 2011; Malmendier, 2021a,b), and consumer demand (Bronnenberg et al., 2012). A pattern emerging in these papers is the persistent effect of experience on high-stakes decisions, showing that people tend to recall experiences that are emotionally salient or contextually similar to the decision at hand.

Our framework assumes that memories stem from one's own past medical experiences (e.g., a previous flu shot) or those shared within personal networks (e.g., an anecdote about a sibling's broken-arm treatment). Upon retrieval, different memories compete with each other; Decisions depend not on the stock of stored experiences but rather on the subset that becomes accessible at decision time. We build our analysis of vaccination decisions on a simple conceptual framework of selective memory inspired by a series of papers by Bordalo et al. (2012, 2022, 2024) and that is reminiscent of the seminal paper on case-based decision theory by Gilboa and Schmeidler (1995).

The decision maker (DM) 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 expected 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

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

Estimates of $\hat{\pi}$ are based, on the one hand, on public information about the true probability of severe adverse events, π, 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 θ as the cost of acquiring information about the true risk of adverse events, though it can more generally be understood as the weight an individual places on personal experience. For example, a lower degree of trust in healthcare authorities would correspond to a higher θ. Equally important, individuals with high health literacy are likely to have better access to information about risks—either because they are more exposed to relevant information or because they can more easily process official statistics. 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 = "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

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

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 DM may use 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 weight–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, σ provides a micro-foundation for how individuals use previous experiences to assess risks; σ 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, σ 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

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

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 θ, and officially provided information π, 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 θ can be interpreted as the cost or difficulty of acquiring information about π, we expect individuals with

low health literacy have higher θ 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 that revolved reimbursement from governmental agents the scandal 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 framework of 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 to the full 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.[3] 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

[3]The measles vaccine is almost always given in combination with the mumps vaccine and the rubella vaccine, combined called the MMR-vaccine. Throughout this paper we'll focus on measles vaccinations, but it will typically coincide with vaccination against also mumps and rubella.

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. Although under-reporting is substantial, all physiological reactions to medication that are secondary to its therapeutic purpose are expected to be reported to the system.

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 **??** 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, the regional Pandemrix dataset covers 13% of Sweden's in 2010.[4]

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 and (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. 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

[4]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.

impossible to isolate the individual cause of the disease; It is likely that all patients in the treatment group 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 rare and even general practitioners may not have been familiar with its specific symptomatology before it became salient in the media.

[Figure 1 about here]

We define a *pre–swine flu* sample of individuals, who are diagnosed with narcolepsy between January 1st 2005 and October 1st 2009–just before the introduction of the vaccination with Pandemrix. 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.[5] These individuals are on average 37 years older than those with Pandemrix-induced narcolepsy.

We contrast medication take-up of narcolepsy patients in the treatment group and the placebo control group to better understand potential differences in clinical manifestations between the two groups. Figure B3 shows the development of prescriptions for common medications used to treat narcolepsy symptoms among the two groups. In line with observations in Gauffin et al. (2022), patients in the post–swine flu sample are in large parts immediately prescribed modafinil, methylphenidate and sodium oxybate. This combination represents the to-date most common pharmaceutical routine for narcolepsy patients. Furthermore, also in line with Gauffin et al. (2022), the prescriptions of dexamphetamine and lisdexamphetamine increased after 2014. Accordingly, the prescribed drug schedules are slightly different in terms of the aforementioned medications for the post–swine flu and the pre–swine flu sample. Changing medication routines takes time, and patients often get started on new prescriptions while fading out the existing medical treatment. For instance, Sodium oxybate, an anesthetic treating cataplexies, did not become available on the market until 2012, implying that patients with narcolepsy-induced Pandemrix were prescribed with it right away while as patients diagnosed before 2010 either stuck to established routines or switched medication later. The differences is also partly explained by the age difference between the two groups. For example, sodium oxybate and atomoxetine are typically not prescribed to older individuals due to elevated risk of respiratory depression. In sum, based

[5]It is estimated that around 4 000 individuals in Sweden suffer from some type of narcolepsy (Gauffin et al., 2022).

on prescription usage, we do not find any clear signs that either of the two groups would have more severe narcolepsy symptoms

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 $i^* = \arg\min_{i \leq N-50}\{ f(i+50) - f(i)\}$. In words, 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 B7 we display the distribution of first vaccinations across time along with defined first date of availability for a sample of 16 region×birth year-cells. In general, the distributions of first vaccinations within these region×birth year-cells are unimodal and concentrated.

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: \textit{neighbors}, defined as individuals living in the same 250×250 m cell as focal sample individuals in 2011; \textit{schoolmates}, defined as individuals attending the same school as the focal individual in 2011; \textit{colleagues}, defined as individuals working at the same plant or establishment in 2021 who are not family members of the treated or control individuals; and \textit{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.

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—which is primarily relevant to the analysis of individuals with narcolepsy—is to isolate the effects from drug-induced adverse events to the effects of having a condition that was not attributed to a drug; We think of *having* a condition such as narcolepsy as a mediating factor of the vaccine hesitancy effect from exposure to adverse events. For example, individuals with narcolepsy may have different experiences with the healthcare system or different social networks that shapes their attitude toward the COVID-19 vaccines. 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 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 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

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

In words, 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

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

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 $\mathbb{E}\left[Y_{ith}|t = \text{Pandemrix}, h = \text{Narc. before}\right]$. Instead, the key assumption is the following

Assumption 1 *Constant narcolepsy effect*

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

That is, the effect on vaccinations of developing narcolepsy is independent of Pandemrix-status. Put differently, the effect of developing narcolepsy before the swine flu pandemic is the same as among Pandemrix-takers.

This assumption allows us to estimate the following object

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

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.

τ_{dd} may be interpreted the causal effect of developing Pandemrix-induced narcolepsy under a conditional independence assumption:

Assumption 2 *Conditional independence*

$$\mathbb{E}\left[Y_i(0)|h_i = \text{Narc.}, X_i\right] = \mathbb{E}\left[Y_i(0)|h_i = \text{No Narc.}, X_i\right] \tag{5}$$

In words, conditional on a set x_i of observable characteristics, developing an adverse event is as good as random. We return to the plausibility of this assumption in Section 5.3

Individuals who developed narcolepsy before the swine flu pandemic may display a different propensity to take the COVID-19 vaccine for two distinct reasons. First, they may believe they are more susceptible to adverse events and, through effects on educational and occupational trajectories, have been placed in social contexts that shape vaccination behavior. Second, they may be more hesitant because they are particularly aware of the health scandal—an awareness that is itself a consequence of the scandal.[6] It is challenging to disentangle the information effect from the narcolepsy effect.

We make progress on identifying which mechanism 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 an estimate for developing narcolepsy prior to the swine flu pandemic (τ_{pre}).

[Figure 2 about here]

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 scandal 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 on vaccination of developing an adverse event such as narcolepsy. 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 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 $Y_i(w, d)$ be the potential outcome

[6] More generally, this violation of SUTVA also applies to individuals without narcolepsy.

under true state $W_i = w, D_i = d$. 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*

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

Among reporters, the effect among developers and non-developers is the same in expectations. This is plausible since individuals, conditional on observable characteristics, do not know if they developed a condition as an adverse event or would have developed it anyways. Put differently, conditional on receiving the cue that a symptom was an adverse event, actually having it makes no difference. Furthermore, we make use of the following assumptions

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

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

Conditional on X_i, among those that develop adverse events, reporting it is independent of future vaccine hesitancy. Critical for our setting, is that both narcolepsy and other recurring adverse events that we consider are severe enough to require healthcare and reported by healthcare professionals. Intuitively, this means that the reporting is outside of the control of the patient while it is also likely that they infer that their condition was an adverse event, independent of whether it was reported. Note that, for narcolepsy we think of almost everyone with vaccine-induced narcolepsy reporting it, but some individuals may at random, not report it. Equivalently, MAR among non-developers states that, among individuals that do not develop an adverse event, reporting is independent of vaccine hesitancy.

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

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

Source irrelevance together with mar(1) directly give us

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

Rare adverse events together with mar(2) give us

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

such that

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

We provide a brief mathematical derivation in Section G In summary, we need four ingredients. First, among individuals who truly experience the adverse event, the decision to report it must be ignorable once we condition on observed covariates; some have doctors report it, whereas others do not. In particular, among those who develop the adverse event, the decision to report it is unrelated to unobserved traits like vaccine hesitancy. 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 the adverse event.

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 (ii) being diagnosed with narcolepsy and reporting narcolepsy as an adverse event. We discuss these concerns in turn.

Characteristics of individuals with narcolepsy diagnoses First, although the pathogenic mechanisms that trigger narcolepsy are still not well understood, the medical literature has found no comorbidities that would systematically predispose particular sub-populations to the disorder. It therefore appears unlikely that underlying biological factors jointly influence both vaccine hesitancy and the risk of developing narcolepsy. Second, 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 swine flu pandemic reduces the chance that cases stem from Pandemrix and instead isolates factors linked to developing and being diagnosed with narcolepsy; it also lets 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.[7] 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 365,000 individuals who developed a heart attack after 2015 along a random sample of individuals residing in Sweden in 2021 information yields a relatively higher AUC of 0.61. Taken together, neither health characteristics or socioeconomic characteristics predicts developing narcolepsy.

Non-random reporting of adverse events Another challenge is potential selection in reporting narcolepsy as an adverse event (ii). 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 that there may be latent vaccine hesitancy, orthogonal to observable socioeconomic and health characteristics, that is correlated with being diagnosed with and reporting narcolepsy as well as 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. In Figure 4 we display differences in the covariates between the diseased individuals and the control individuals, conditioning only on birth year. Individuals in the treatment group have higher income and years of schooling. There is no clear pattern regarding their health status, on the

[7]Specifically, each treated unit is matched to 100 untreated units with the same birth year.

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 likely an artifact of the fact that we do observe control individuals from the larger metropolitan areas. Individuals born before or in 1990 (panel b) report narcolepsy as an adverse event somewhat later than individuals born after 1990 (panel a) but appear to not be more selected in terms of socioeconomic characteristics, despite the fact that these are individuals who were more likely to have developed narcolepsy anyways.

In Figure B5 we display the same coefficients for the pre-sample. Recall that, for the pre-sample, we compare individuals who were diagnosed with narcolepsy prior to the swine flu pandemic 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 4. 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 (placebo control 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 computing propensity scores with gradient boosting and coarsened exact matching (CEM) instead of propensity score matching. The former aims at better handling interactions, accommodating functional forms that are not well captured by the logit model. A 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. The latter aims at dealing with the fact that we have few treated units, making the maximum-likelihood estimates in the logit propensity-score model unstable. Instead, matching is done exactly on gender, birthyear and coarsened versions of the continuous covariates: Parents' years of schooling, parents' income, parents' number of healthcare visits, parents' number of drug prescriptions.

[Figure 4 about here]

In Figure B6 we display how the different methods perform in terms of achieving balance across the covariates. The differences in performance are small. If anything, CEM performs somewhat worse than the NN and gradient boosting propensity score matching in achieving balance.

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-samples, 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, a notion we explore using large-scale descriptive evidence on the role of prior experiences during the COVID-19 pandemic. In Figure 5 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 heath related variables that we observe, namely their diagnoses, medical drugs and reported side effects between 2010 and 2020.Our preferred method for computing propensity scores 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 5 about here]

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 access data on genes and include vaccination status of family members. Similar to us, they find that, apart from income, medication history has the highest predictive value. The prediction exercise provides descriptive evidence that previous experiences with healthcare at large shape attitudes toward vaccines. Taken together, while socioeconomic characteristics strongly predicts vaccination uptake, the incremental increase from a decade of personal health experiences 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 6. 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. We report results for τ_{post}, the mean difference between treated individuals (post—swine flu narcolepsy) and matched individuals that received Pandemrix; τ_{pre}, the mean difference between placebo control individuals (pre—swine flu narcolepsy) and matched individuals from the general Swedish population; and τ_{dd}, the difference between τ_{post} and τ_{pre}.

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 were vaccinated. We also find a lower vaccination rate for the placebo control group of individuals that developed narcolepsy before vaccination with Pandemrix. As discussed in Section 5, it is unlikely that there is anything inherent about individuals with narcolepsy. Instead, the plausible explanation is that these individuals have more information about and are more aware of the medical scandal and therefore more hesitant. Combining τ_{post} and τ_{pre} yields a large effect of 35 percentage points lower vaccination rate. Given the informational spillovers, we interpret this estimate as a lower bound of the effect of developing Pandemrix-induced narcolepsy.

To gauge the economic significance, we compare the magnitude of the effects to the socioeconomic gradient in vaccination rate which is displayed in Figure B1. There is a meaningful gradient in immunization outcomes along socioeconomic characteristics but 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 higher than five percentage points. These differences are large given the high baseline vaccination rate, but small relative to the effect sizes of being exposed to severe adverse events.

We next consider the number of doses, conditional on getting one dose, to account for vaccine hesitancy expressed in incomplete vaccination schedules (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 seem to be less hesitant to also take additional doses required to meet the full vaccine schedule.

Considering the number of days elapsed until the first vaccination, diseased individuals delay their first vaccination by approximately 40 days, family members by approximately a week, 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 to see if others experience adverse events from the COVID-19 vaccine. Close family members delay their vaccination with 10 days—similar to the difference between individuals with only primary education (nine years of schooling) and individuals with post tertiary degree.

We replicate the results in Figure 6 using coarsened exact matching and propensity scores computed using gradient boosting in Figure E2 and Figure E1. Our results remain virtually identical with the exception of *Number of Doses* where we find negligible effects for CEM. In Figure B8 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 reweigh 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 extrapolate the effect by fitting a spline of vaccination uptake against age in the pre–swine flu sample and applying it to the age distribution of post–swine flu sample. None of these two exercises changes the estimate of τ_{pre} meaningfully. This suggests that our estimated effect of having narcolepsy in the placebo control group reflects the counterfactual effect of developing narcolepsy, that was not Pandemrix-induced, at the age of when individuals developed Pandemrix-induced narcolepsy.

[Figure 6 about here]

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 will 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, regular 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 in closely related family of the patients makes the experience particularly easy for them to recall and use for simulation.

In Table B2 we show results where the treated individuals instead consist of individuals that developed mild adverse events from Pandemrix. These adverse events are mostly self-reported and examples include fever, connective tissue pain and headache. The coefficients are much smaller than in Figure 6 but still meaningful. The fact that they developed mild symptoms is unlikely to change their belief about the risk of severe adverse events in itself. Instead, we hypothesize that they more easily recall the experience of developing narcolepsy from Pandemrix and therefore become more vaccine hesitant.

Learning about predisposition The finding that direct family members have similar but weaker effects on immunization outcomes as illustrated in Figure 6 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 B9, we show results for the partners of the diseased individuals. For partners we find large estimates within the size range of those found for diseased individuals and close family members. Keeping in mind the endogeneity in partner choice, we interpret these findings as evidence that the results in Figure 6 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.

Mechanisms alternative to altered perceived costs Our simple conceptual framework suggests that treated individuals are primarily influenced by altered beliefs about the risk of adverse events. The framework also allows for the theoretical possibility that they respond to changes in perceived infection risk, or, put differently, to altered beliefs about the benefits of vaccination. This interpretation fits with the idea that they shy away from the healthcare system altogether. To explore the possibility that decreased expected benefits of vaccination

explain the worsened immunization outcomes, we elicit the perceived benefits of vaccination. We assess treated individuals' concern about the disease through two different 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 individuals who received Pandemrix.

The results are displayed in Figure 9. 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 be partly explained by 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 story that we do not rule out, consistent with these findings, is that treated individuals do not expose themselves as much to other individuals, thereby reducing their risk of contracting COVID-19.

Experiencing the benefits of vaccination We augment the analysis by turning to individuals who experience the benefits of vaccination, rather than its costs. To do this, we consider individuals are develop diseases, that vary in the degree of similarity to COVID-19. Point being that developing potentially vaccine-preventable conditions are more aware of the consequences of foregoing vaccination. We focus on individuals that develop conditions with similar symptoms to COVID-19 and that are (i) vaccine-preventable, (ii) infectious but not vaccine-preventable. In Figure B10, we consider the same outcome variables as previously, but where the independent variable now takes the value 1 if an individual has been hospitalized for influenza at some point between 2015 and 2020, right before vaccination against COVID-19 started. Controls are nearest neighbors matched on socioeconomic and health characteristics, with exact matching on gender and birth year. Each regression is reweighted to match a common age distribution. Our results provide evidence that the exposure to the benefits of vaccination reduces vaccine hesitancy. In particular, individuals who developed COVID-19 during the initial stages of the pandemic have significantly improved immunization outcomes. Qualitatively similar but less pronounced results holds for influenza that is also vaccine-preventable and infectious and Upper respiratory tract infections that is 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 B12 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 in a way such 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 8 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 a salient association for these individuals when deciding whether to take the COVID-19 vaccine, yet they appear unaffected. Our interpretation is that the frequency of exposure in itself fundamentally shapes the salience and ease of recall of the narcolepsy episode, and thereby its influence on subsequent vaccination decisions.

[Figure 8 about here]

Attributing other diseases to Pandemrix As stated in the data section, the best available evidence suggests that Pandemrix did not cause other diseases than narcolepsy. However, 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 B13 we display COVID-19 vaccination rates as a function of the year in which it was first diagnosed for these different diseases. Since these diseases are much more common than narcolepsy. 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 B14 we display the yearly number of specialist healthcare visits across time. In a short period after developing narcolepsy patients make more visits to medical professional, likely in examination phase, before they are diagnosed with narcolepsy, after that their 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 exposure 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 the Pandemrix–narcolepsy episode. Yet, the broad take-away is that individuals put a lot of weight on own personal experiences, in particular when faced with novel risks, a central question is why they do so. 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.

[Figure 7 about here]

In Figure 7, panel A we examine the role of the size of the database. We consider 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, (ii) The number of healthcare visits.

We find that Individuals who have taken a greater number of unique drugs react some what 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 (θ 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 estimates 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). For individuals with negative experiences, experiences crowd out officially provided information, such that the level of health literacy does not matter. 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 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 now, 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 remarkable similarities with the deployment of the COVID-19 vaccine. Moreover, we have strong reasons to believe that developing and reporting narcolepsy as an adverse event is plausibly 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 consider the effects of severe adverse events from all pharmaceutical products (including all vaccines apart from pandemrix).

Unlike the rare, vaccine-specific 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 national register captures a broad and 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 reporting vaccine adverse events and 33 000 individuals reporting non-vaccine adverse events between 2014 and 2020—right before vaccination against COVID-19 started in Sweden. We focus on the effect of *severe* adverse events that are reported by healthcare professionals. As displayed in Figure 10, 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 reveal a greater number of sick leave days, healthcare visits and drug prescriptions, but do not appear positively selected on socioeconomic characteristics.[8] To maximize the number of observations we impose no restriction on the individuals in the sample having taken pandemrix. This, however, means that we can directly assess how individuals reporting adverse events differ in their propensity to vaccinate by considering their vaccination rate with pandemrix. Individuals that later reported adverse events are *more* likely to have taken the Pandemrix vaccine. This is strong evidence that our results—both on pandemrix-induced narcolepsy and on non-vaccine adverse events that we are providing in this section—can not be explained by latent vaccine hesitancy that correlates with both the propensity to report adverse events and willingness to take the COVID-19 vaccine.

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

[Figure 10 about here]

The individuals that report general adverse events are relatively older than those that developed Pandemrix-

[8]Individuals that report mild adverse events are, however, positively selected in terms of socioeconomic characteristics.

induced narcolepsy. In 2021, when the vaccine was rolled out, the median age of the individuals that reported adverse events was 63. At this age, forgoing or delaying COVID-19 vaccination entails substantial health risks compared to the average individual that developed narcolepsy.

What does experiencing an adverse event entail? Figure 11 shows health outcomes around the time of reporting an adverse event 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 beyond 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 any differences in effects between vaccine adverse events and non-vaccine adverse events do not 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.

[Table 1 about here]

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. Interestingly, the coefficients are roughly proportional to the results that we find for narcolepsy with the coefficients for vaccine taken and for days elapsed being about 1/20 for non-vaccine adverse events and about 1/3 for vaccine adverse events of the coefficients that we find for narcolepsy. For family members we find precisely estimated null results 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 B3 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.

Stickiness of experiences In Figure B11 we break up the results by year of reporting adverse events between 2013 and 2020. There is no indication that the effect decays over time, highlighting that, while the adverse

events are not on top of individuals' minds, the memories are reactivated years later when facing the decision to vaccinate against COVID-19. A plausible concern is that there is an interaction between the salience of adverse events and developing an adverse event. In other words, individuals that developed adverse events shortly after the narcolepsy scandal unfolded are more likely to be affected. The salience interaction could explain the dip in uptake among those who experienced vaccine adverse events in 2016, when public debate over narcolepsy compensation resurfaced as the government assumed responsibility for compensating 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 position 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. In particular, we focus on children's vaccines against measles, mumps and rubella (MMR), diphtheria and pneumococcal disease. 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. 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. In consequence, easy access to scientific evidence lowers information costs and increases reliance on external sources. 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 θ 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 their childhood.

We define a sample of individuals born in 2013 and onward who are registered as living 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 4.1% of children, parents abstain from having the measles vaccine administered, compared to 1.3 % for diphtheria and 2 % for pneumococcus. The risk of developing any of these diseases for unvaccinated individuals is very low. Instead, a plausible explanation is the perceived link between measles and autism that we will investigate later in this section. Panel B of Table 2 illustrates that children vaccination decisions correlate highly. This observation suggests that parental vaccine hesitancy reflects a general sentiment rather than concerns about a specific vaccine.

[Table 2 about here]

9.1 Exposure to Narcolepsy

In a first step, we study effects among children as well as nieces and nephews of individuals that themselves developed narcolepsy. The dependent variable equals one if a child completed 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. Table 3 displays results for children in panel A and nieces and nephews in panel B of individuals that developed narcolepsy and matched control individuals. Column 5 contains the estimates for an aggregated index of adherence, defined as the average vaccine adherence for the three different diseases examined. Given our data limitations, we are restricted to a few individuals in the treatment group who have children born after 2013, which limits the precision of our estimates.Since we only have about 65 parents who had children by the end of the observation period, We deploy a simple matching strategy where we match exactly on coarsened versions of socioeconomic characteristics of the parents as well as birth month. The data does not allow to observe if the children get the vaccine later than what is scheduled. Furthermore, we focus on the first difference τ_{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 intervals computed using and confidence interval obtained based on Fisher's exact test. This avoids reliance on asymptotic normality for the odds-ratio. Against this background, the negative point estimates are large, indicating that individuals who develop narcolepsy also abstain from fully vaccinating their children with well established vaccines.

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.

[Table 3 about here]

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 any 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 of adverse events across children within the same family. To this end, we estimate the following regression equation:

$$y_{pb} = \delta_b + \gamma_p + \beta T_{pb} + \epsilon_{pb} \tag{6}$$

Where δ_b denote birth order fixed effects, γ_p denote parent fixed effects and T_{pb} is equal to one for a child with birth order b is born to a parent that reports a severe adverse event, after the parent has reported an adverse event. y_{pb} is the vaccination status for the child. We construct separate sibling groups based on each parent and assign weight 0.5 for each child that appears 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.

[Table 4 about here]

For non-vaccine adverse events we find no negative effects on vaccine 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 vaccinations we find 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 does cause individuals to be more reluctant vaccinating their children with well-established and safe vaccines.

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

Finally, we aim to understand whether adverse events need to be backed up by scientific evidence to produce the substantial negative effects on immunization outcomes observed in the previous sections. In particular, we isolate a scenario where an adverse event is likely perceived to result from medical treatment but lacks scientific recognition. To this end, we exploit an infamous scandal surrounding the MMR vaccine and study younger siblings of children who developed autism soon after receiving the measles vaccine.

In 1998, a controversial study of 12 children was published in The Lancet, suggesting a potential link between the measles, mumps, and rubella (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. Nevertheless, as is indicated by the lower vaccination rates for the MMR vaccine compared to the other vaccines in table 2, the publication had a lasting impact on public perceptions of vaccine safety.[9]

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 12. 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% lower take-up relative to the control mean of 90%), 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 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 controversial 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 severe 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. In particular, 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 exposure to the adverse event of narcolepsy leads

[9]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.

to significantly lower COVID-19 vaccine take up, uptake of fewer doses, and delay in vaccination for the vaccinated diseased 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 reoccurring 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 suggest 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. Given this, a natural question is what can restore actually trust and improve immunization outcomes. The role of reimbursing individuals that have had negative experiences has not been studied. Another interesting venue is the role of positive experiences 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, A. and J. Spiess (2022). Robust post-matching inference. *Journal of the American Statistical Association 117*(538), 983–995.

Adhikary, D., S. Barman, R. Ranjan, and H. Stone (2022). A systematic review of major cardiovascular risk factors: a growing global health concern. *Cureus 14*(10).

Ahmed, S. S., P. H. Schur, N. E. MacDonald, and L. 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.

Alsan, M. and M. Wanamaker (2017, August). Tuskegee and the Health of Black Men*. *The Quarterly Journal of Economics 133*(1), 407–455.

Anderberg, D., A. Chevalier, and J. 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, B. and F. Annan (2023, September). 'we are not guinea pigs': The effects of disclosure of medical misconduct on vaccine compliance. Technical report.

Ashraf, N., G. Bryan, A. Delfino, E. Holmes, L. 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 . . . .

Barateau, L., R. Lopez, and Y. Dauvilliers (2016). Treatment options for narcolepsy. *CNS drugs 30*(5), 369–379.

Baumann, C. R., E. Mignot, G. J. Lammers, S. Overeem, I. Arnulf, D. Rye, Y. Dauvilliers, M. Honda, J. A. Owens, G. Plazzi, et al. (2014). Challenges in diagnosing narcolepsy without cataplexy: a consensus statement. *Sleep 37*(6), 1035–1042.

Bordalo, P., G. Burro, K. Coffman, N. Gennaioli, and A. Shleifer (2024, 06). Imagining the future: Memory, simulation, and beliefs. *The Review of Economic Studies*, rdae070.

Bordalo, P., J. J. Conlon, N. Gennaioli, S. Y. Kwon, and A. Shleifer (2022, 08). Memory and probability*. *The Quarterly Journal of Economics 138*(1), 265–311.

Bordalo, P., N. Gennaioli, and A. Shleifer (2012, 06). Salience theory of choice under risk. *The Quarterly Journal of Economics 127*(3), 1243–1285.

Bronnenberg, B. J., J.-P. H. Dubé, and M. Gentzkow (2012, May). The evolution of brand preferences: Evidence from consumer migration. *American Economic Review 102*(6), 2472–2508.

Chang, L. V. (2018). Information, education, and health behaviors: Evidence from the mmr vaccine autism controversy. *Health economics 27*(7), 1043–1062.

Deer, B. (2020). *The doctor who fooled the world: Science, deception, and the war on vaccines*. JHU Press.

Edwards, K., G. Hanquet, S. Black, E. Mignot, C. Jankosky, T. Shimabukuro, E. Miller, H. Nohynek, and P. Neels (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.

Folkhälsomyndigheten (2024). Barnvaccinations programmet i sverige 2024. Technical report.

Gauffin, H., I. Boström, S. G. Berntsson, A. Kristoffersson, M. Fredrikson, and A.-M. Landtblom (2024). Characterization of the increase in narcolepsy following the 2009 h1n1 pandemic in sweden. *Journal of Clinical Medicine 13*(3), 652.

Gauffin, H., T. Fast, A. Komkova, S. Berntsson, I. Boström, and A.-M. Landtblom (2022). Narcolepsy treatment in sweden: An observational study. *Acta Neurologica Scandinavica 145*(2), 185–192.

Gilboa, I. and D. Schmeidler (1995). Case-based decision theory. *The Quarterly Journal of Economics 110*(3), 605–639.

Giulietti, C., M. Vlassopoulos, and Y. Zenou (2023). When reality bites: Local deaths and vaccine take-up. *European Economic Review 156*, 104463.

Graeber, T., C. Roth, and F. Zimmermann (2024, 06). Stories, statistics, and memory*. *The Quarterly Journal of Economics 139*(4), 2181–2225.

Hallberg, P., H. Smedje, N. Eriksson, H. Kohnke, M. Daniilidou, I. Öhman, Q.-Y. Yue, M. Cavalli, C. Wadelius, P. K. Magnusson, et al. (2019). Pandemrix-induced narcolepsy is associated with genes related to immunity and neuronal survival. *EBioMedicine 40*, 595–604.

Hartonen, T., B. Jermy, H. Sõnajalg, P. Vartiainen, K. Krebs, A. Vabalas, A. Metspalu, T. Esko, M. Nelis, G. Hudjashov, T. Leino, H. Nohynek, J. Sivelä, R. Mägi, M. Daly, H. M. Ollila, L. Milani, M. Perola, S. Ripatti, A. Ganna, FinnGen, and E. B. R. Team (2023). Nationwide health, socio-economic and genetic predictors of covid-19 vaccination status in finland. *Nature Human Behaviour 7*(7), 1069–1083.

Holmberg, S. and P. Hedberg (2020). Vilka vaccinerade sig mot svininfluensan 2009/10? Technical report, The SOM Institute.

Holmes, E. A., A. Mathews, T. Dalgleish, and B. Mackintosh (2006). Positive interpretation training: Effects of mental imagery versus verbal training on positive mood. *Behavior Therapy 37*(3), 237–247.

Insani, W. N., C. Whittlesea, H. Alwafi, K. K. C. Man, S. Chapman, and L. Wei (2021). Prevalence of adverse drug reactions in the primary care setting: A systematic review and meta-analysis. *PLOS ONE 16*(5), e0252161. eCollection 2021.

Komagamine, J. (2024, mar). Prevalence of urgent hospitalizations caused by adverse drug reactions: a cross-sectional study. *Scientific Reports 14*(1), 6058.

LaBar, K. S. and R. Cabeza (2006). Cognitive neuroscience of emotional memory. *Nature Reviews Neuroscience 7*(1), 54–64.

Larson, H. J., E. Gakidou, and C. J. Murray (2022). The vaccine-hesitant moment. *New England Journal of Medicine 387*(1), 58–65.

Larson, H. J., C. Jarrett, E. Eckersberger, D. M. Smith, and P. 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, S. and E. Montero (2021, April). The legacy of colonial medicine in central africa. *American Economic Review 111*(4), 1284–1314.

Mahlios, J., A. K. De la Herrán-Arita, and E. Mignot (2013). The autoimmune basis of narcolepsy. *Current Opinion in Neurobiology 23*(5), 767–773. Circadian rhythm and sleep.

Malmendier, U. (2021a). Experience effects in finance: Foundations, applications, and future directions. *Review of Finance 25*(5), 1339–1363.

Malmendier, U. (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, U. and S. Nagel (2011). Depression babies: do macroeconomic experiences affect risk taking? *The quarterly journal of economics 126*(1), 373–416.

Malmendier, U., S. Nagel, and Z. Yan (2021). The making of hawks and doves. *Journal of Monetary Economics 117*, 19–42.

Martinez-Bravo, M. and A. Stegmann (2021, May). 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, A. Nicoll, F. Dapos;Ancona, P. L. Lopalco, K. Johansen, A. M. Wasley, P. Jorgensen, D. Lévy-Bruhl, C. Giambi, P. Stefanoff, L. Dematte, D. O'Flanagan, and C. the VENICE project gatekeepers group (2012). Influenza a(h1n1)pdm09 vaccination policies and coverage in europe. *Eurosurveillance 17*(4).

Motta, M. and D. 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, J., F. Vallières, R. P. Bentall, M. Shevlin, O. McBride, T. K. Hartman, R. McKay, K. Bennett, L. Mason, J. Gibson-Miller, et al. (2021). Psychological characteristics associated with covid-19 vaccine hesitancy and resistance in ireland and the united kingdom. *Nature communications 12*(1), 29.

Oster, E. (2018, January). Does disease cause vaccination? disease outbreaks and vaccination response. *Journal of Health Economics 57*, 90–101.

Partinen, M., B. R. Kornum, G. Plazzi, P. Jennum, I. Julkunen, and O. 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.

Shattock, A. J., H. C. Johnson, S. Y. Sim, A. Carter, P. Lambach, and R. C. Hutubessy (2024, May). Contribution of vaccination to improved survival and health: modelling 50 years of the expanded programme on immunization. *The Lancet 403*(10441), P2307–2316.

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].

12 Tables and Figures

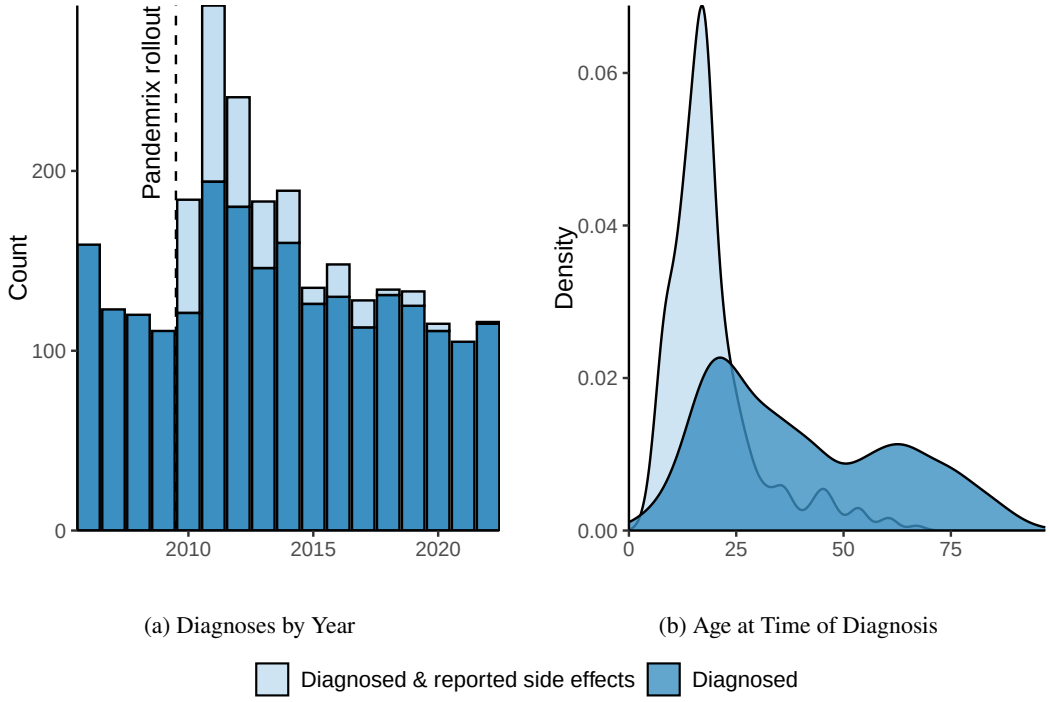

(a) Diagnoses by Year

(b) Age at Time of Diagnosis

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 a side effect to the Medical Products Agency in dark blue and those that were diagnosed but did not report it as a side effect 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.

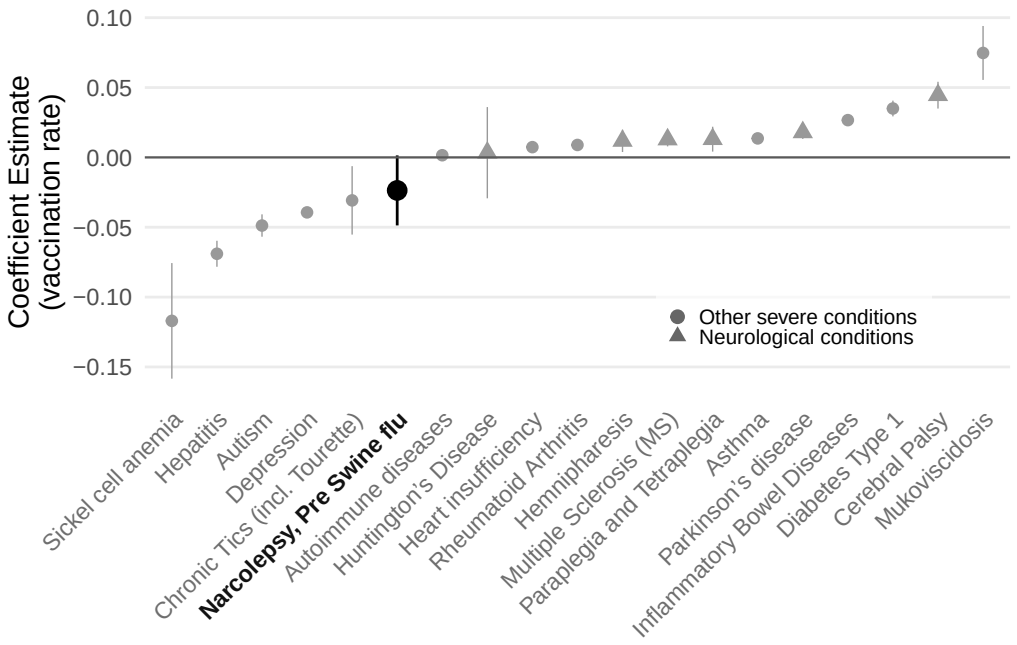

Figure 2: Saliency of Pandemrix Induced Narcolepsy.

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.

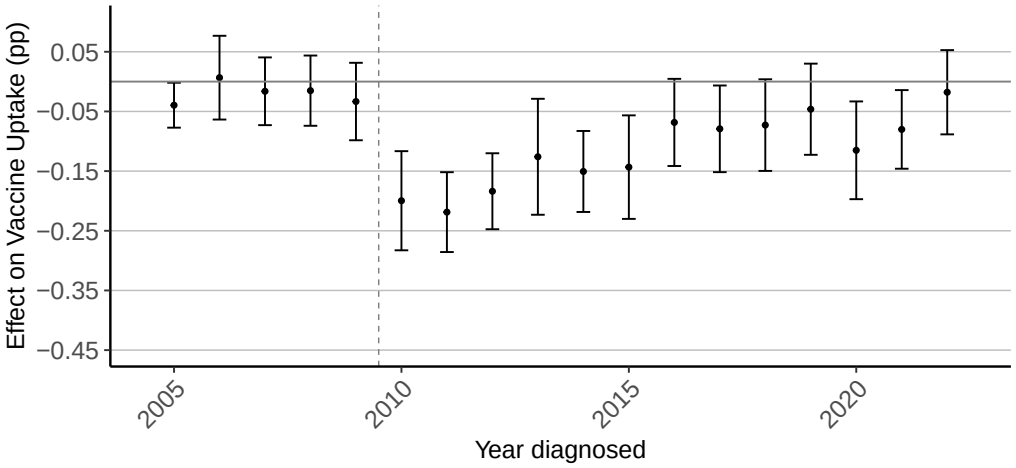

Figure 3: Effect on COVID-19 vaccine uptake of Narcolepsy Diagnoses Across Time

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.

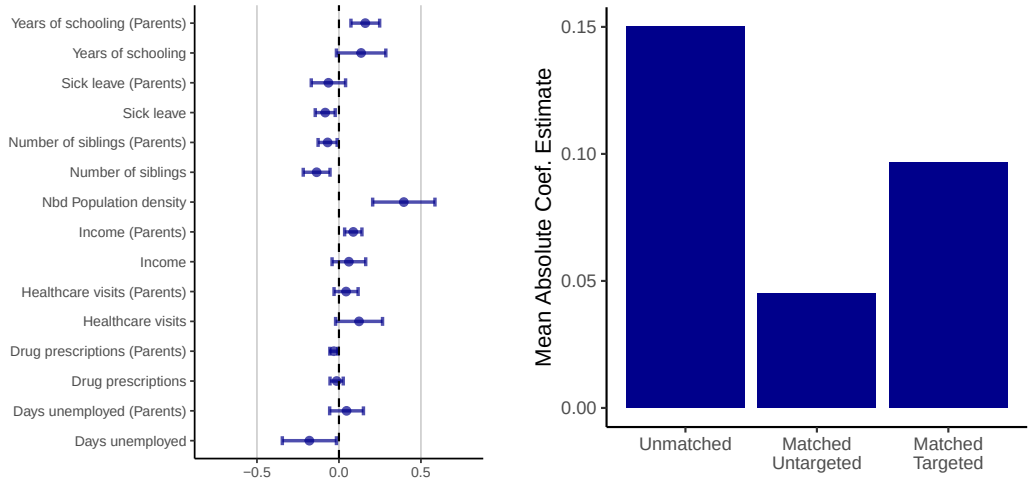

(a) Coefficient Forest Plot

(b) Aggregate Imbalance before/after matching

Figure 4: 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 columns 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.

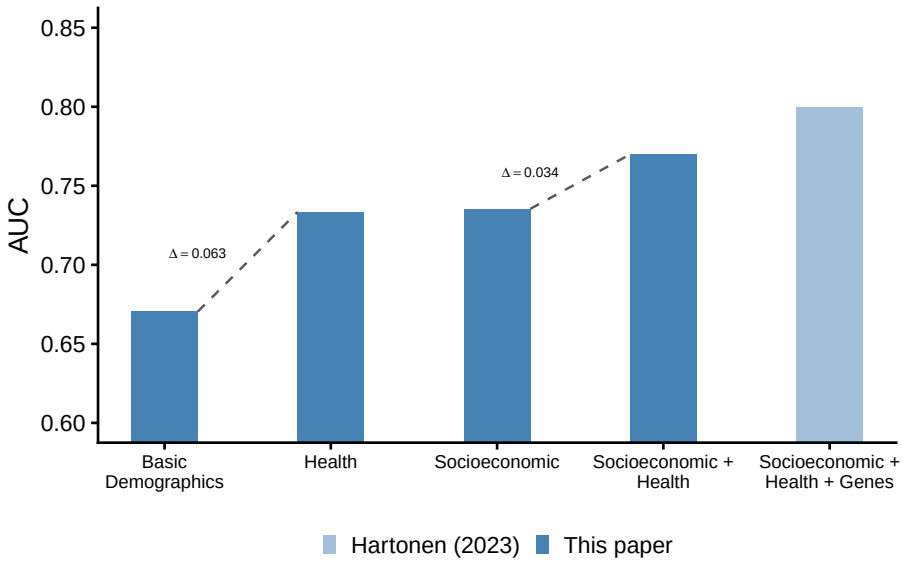

Figure 5: Predicting Who Takes the COVID-19 Vaccine

Notes: Out of Sample AUC-values for predicted probabilities using a neural network. Sample restricted to individuals between 40 and 60 years old who do not belong to COVID-19 risk groups. *Demographic variables* include number of siblings, gender, birth year, country origin and gender. **Socioeconomic characteristics** include income, years of schooling, days unemployed and days sick and demographic characteristics*Health characteristics* include the history of diagnoses, drugs, reported adverse events and demographic characteristics. All models are evaluated out-of-sample. Models are described in details in F

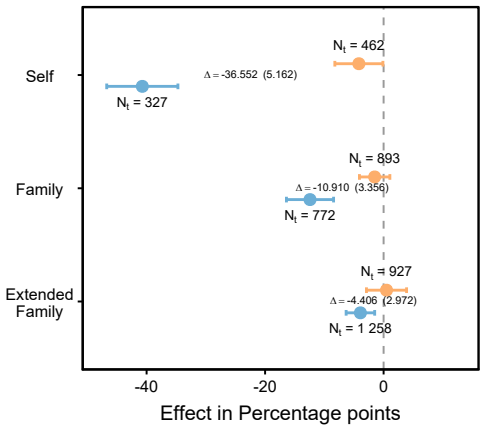

(a) Vaccine Taken

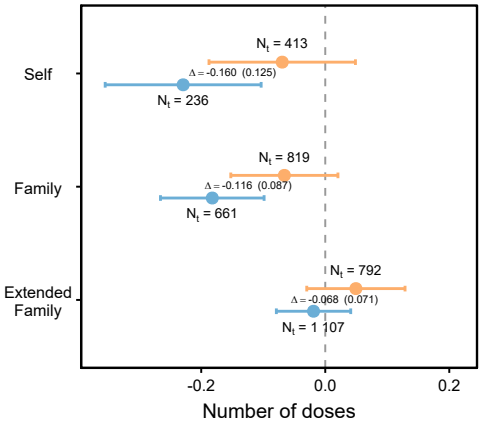

(b) Number of Doses

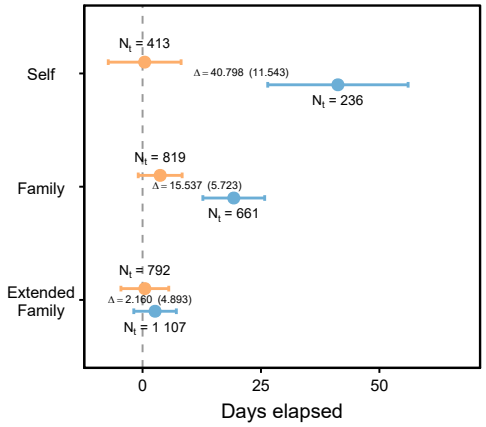

(c) Days until First Dose

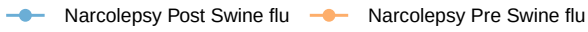

Figure 6: Main Results – Diseased Individuals and Family Members.

Notes: This figure displays coefficients corresponding to τ_{post} and τ_{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, i.e. by the treated or control individual that a network member is related to.

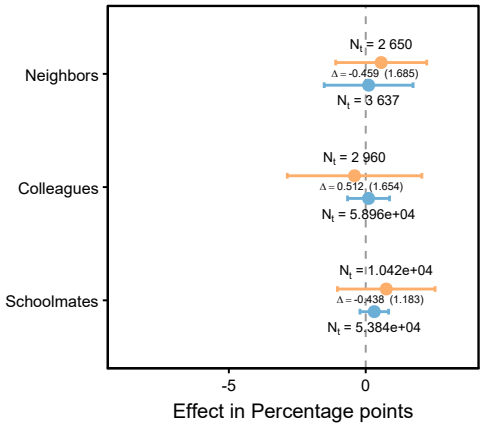

(a) Vaccine Taken

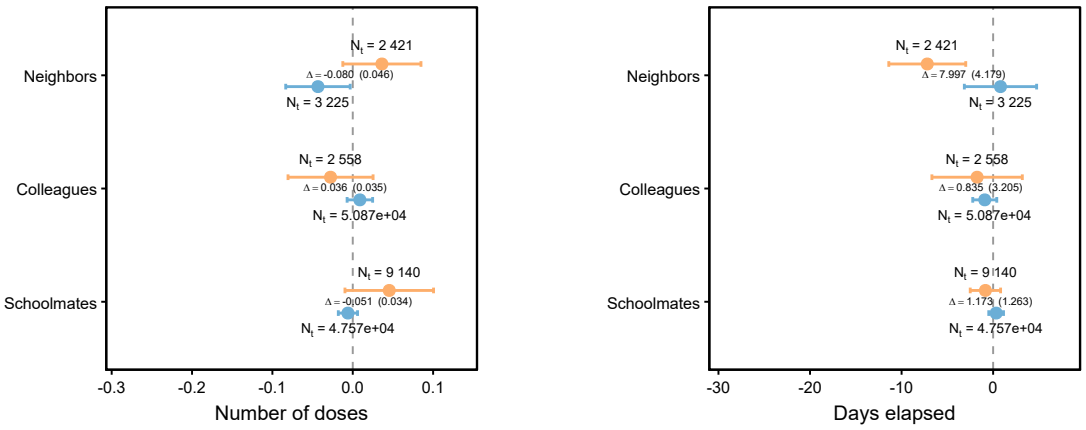

(b) Number of Doses

(c) Days until First Dose

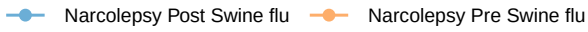

Figure 8: Main Results – Extended networks.

Notes: This figure displays coefficient corresponding to τ_{post} and τ_{pre} in eq. (4). The first row corresponds to coefficients for colleagues of core sample individuals in 2021. The second row corresponds to coefficients schoolmates to the sample individuals in 2011. The third row corresponds to coefficient for individuals that lived in the same neighborhood 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.

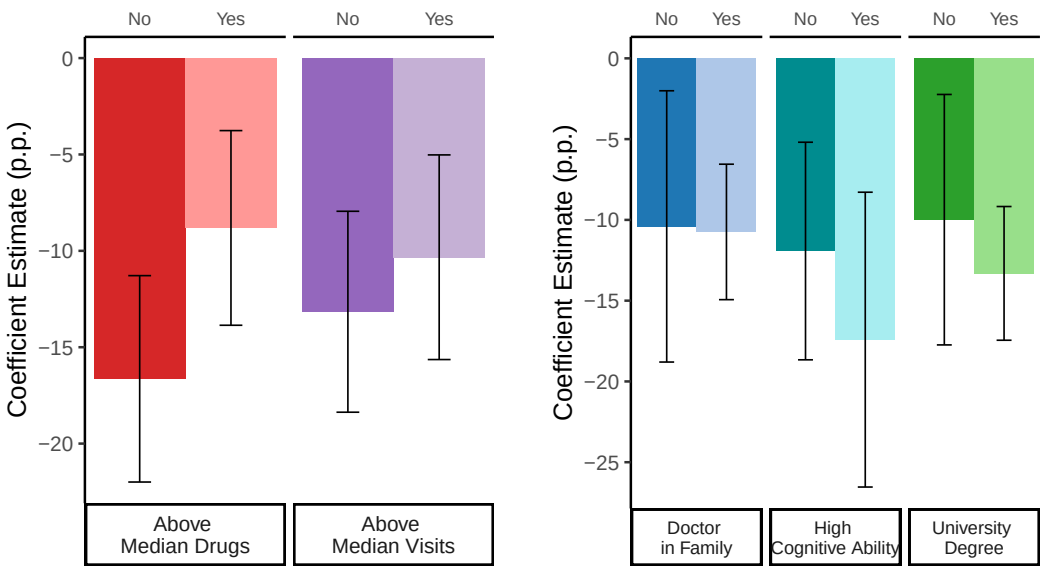

(a) Size of Database

(b) Health Literacy

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

Notes: This figure show 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 θ. We restrict the attention to τ_{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 ≈6) between 2005 and the point of developing side effect relative to its birth year peers. *Above Median Visits* is defined equivalently but for visits to specialized healthcare between 2005 and 2009 (≈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.

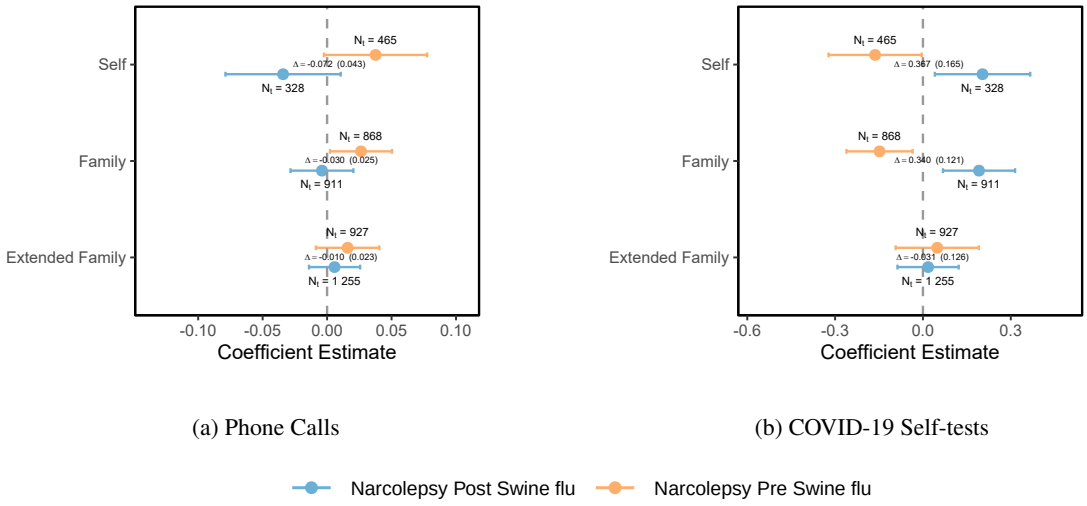

(a) Phone Calls

(b) COVID-19 Self-tests

Figure 9: Eliciting the Perceived Benefits of COVID-19 vaccination

Notes: This figure displays coefficients corresponding to τ_{post} and τ_{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.

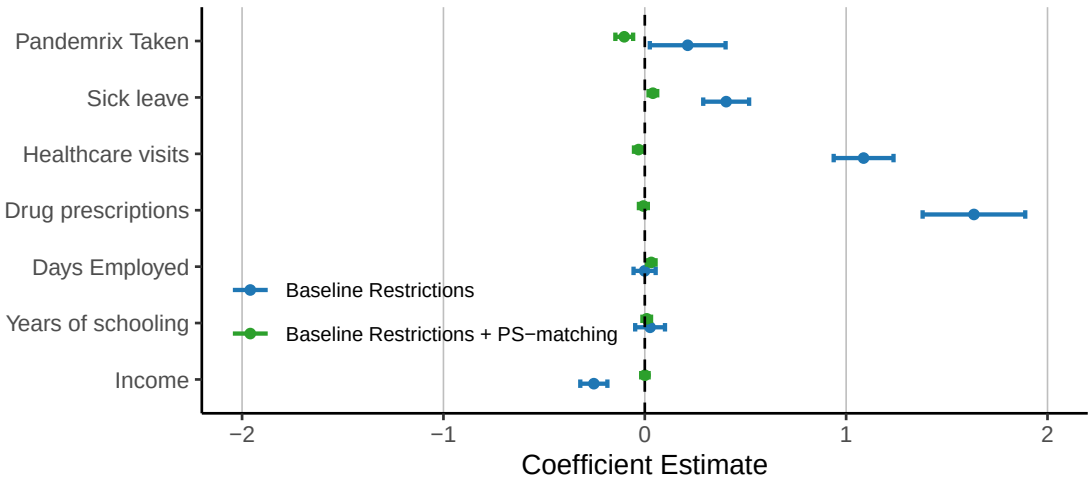

Figure 10: Balance in Pre-treatment Characteristics – General Adverse Events

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.

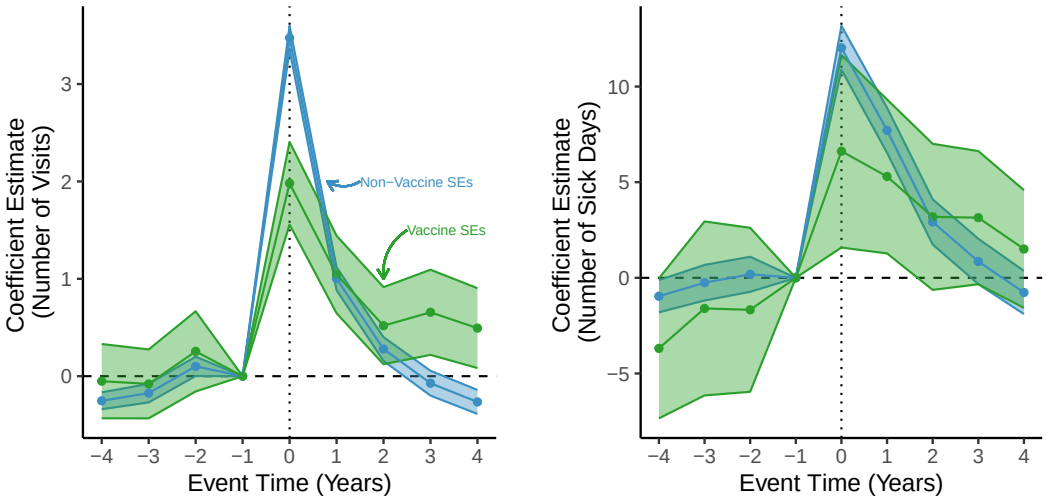

(a) Number of Healthcare Visits

(b) Number of Sick Days

Figure 11: DiD-Estimates of Health Effects from Adverse Events

Notes: The figure shows estimated health outcomes around the reporting of a severe side effect relative to a matched control group for vaccine side effects and non-vaccine side effects. **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 calender year fixed effects.

Table 1: Side Effects and Immunization Outcomes

	Non-Vaccine Adverse Events			Vaccine Adverse Events		
	Vaccine Taken	Number of Doses	Days Elapsed	Vaccine Taken	Number of Doses	Days Elapsed
Diseased Individuals						
Side Effect	-0.006	-0.038	-3.19	-0.075	-0.056	10.3
	(0.003)	(0.015)	(0.68)	(0.019)	(0.073)	(3.51)
N. Treated	12 230	11 326	11 326	672	553	553
Family Members						
Side Effect	0.002	0.030	0.30	-0.037	-0.092	5.91
	(0.002)	(0.010)	(0.42)	(0.011)	(0.040)	(1.92)
N. Treated	28 509	26 323	26 323	1 643	1 431	1 431
Children						
Side Effect	0.009	0.021	-0.74	-0.042	0.056	7.42
	(0.004)	(0.009)	(0.56)	(0.022)	(0.050)	(3.55)
N. Treated	24 188	19 670	19 670	669	513	513
Partner						
Side Effect	-0.001	-0.008	-1.32	-0.009	-0.056	7.90
	(0.003)	(0.020)	(0.79)	(0.015)	(0.096)	(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 a side effect between 2015 and 2020. Columns 1–3 use an indicator for reporting any side effect between 2015 and 2020 as the regressor ("General side effects"). Columns 4–6 instead use an indicator for reporting a vaccine side effect ("Vaccine side effects"). 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.

Table 2: Childhood Vaccines – Descriptive Statistics

Panel A: Share Adhering

	Diphtheria	Pneumococcus	Measles	Covid (Parents)
Abstention	5.6	7	10.0	11.0
Partial Adherence	4.5	8		78.9
Full Adherence	89.8	85	90.0	0.0

Panel B: Correlation matrix, Partial Adherence

	Diphtheria	Pneumococcus	Measles	Covid (Parents)
Diphtheria	1.000			
Pneumococcus	0.922	1.000		
Measles	0.706	0.686	1.000	
Covid (Parents)	0.100	0.107	0.124	1.000

Notes: This figure displays the uptake of vaccines in the Swedish children vaccination program. **Panel A:** Abstention is defined as having received zero doses of the vaccine. Partial adherence defined as having received one or two doses for diphtheria and pneumococcus. Full adherence defined as having received three or more doses for diphtheria and for pneumococcus and having received one or more dose for measles. **Panel B:** Correlations in having received at least one dose. COVID-19 (parents) is defined as the average of parental COVID-19 vaccination status.

Table 3: Exposure to Narcolepsy and Children's Vaccine Outcomes

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

Table 4: General Side Effects and Children's Vaccines

	Non-vaccine Adverse Events				Vaccine Adverse Events			
	Measles	Diphtheria	Pneumococcal	Index	Measles	Diphtheria	Pneumococcal	Index
Treated	0.009	0.018	0.006	0.011	-0.102	-0.028	-0.100	-0.077
	(0.016)	(0.017)	(0.019)	(0.015)	(0.050)	(0.068)	(0.063)	(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 β 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 side effects refers to side effects reported from vaccines (ICD-10-SE J07), excluding adverse events from Pandemrix and COVID-19 vaccines. General side effects refer to side effects reported from other types of drugs. Standard errors are clustered by parent and child.

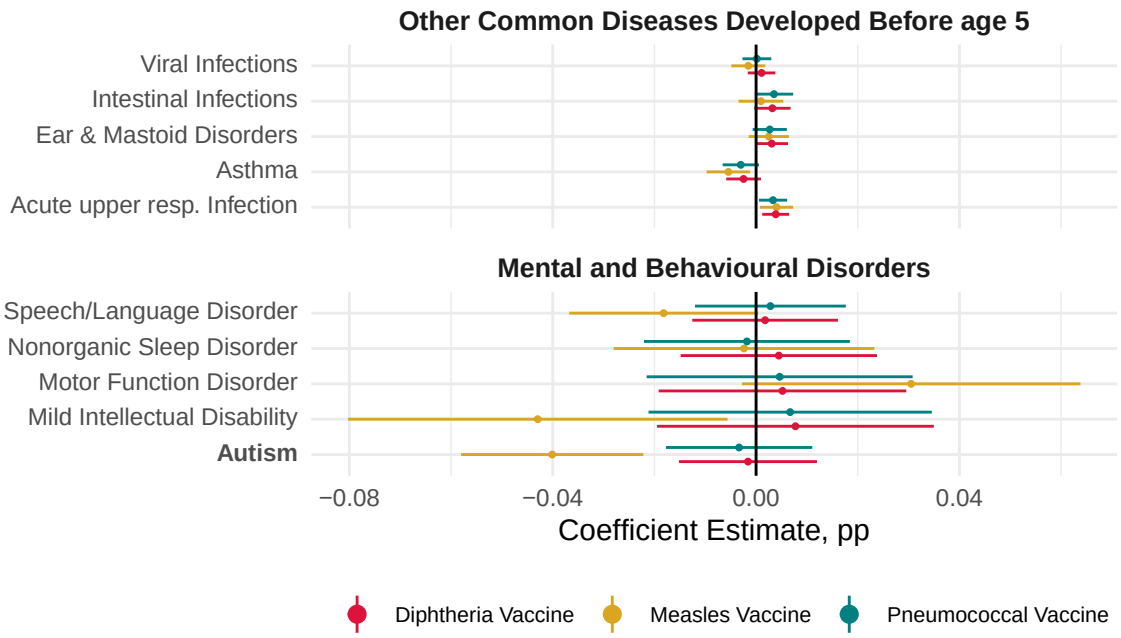

Figure 12: Vaccination Uptake After an Older Sibling Is Diagnosed With Autism

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 vaccinates, is 6,347. Standard errors are clustered by parent.

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 heard 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 side effects 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 form several countries to supply a vaccine against the swine flu. Sweden had already a 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 side effects.

Appendix B. Additional Descriptive Statistics & Results

Table B1: Predictions

	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 probabilties. Probabilities are computed using XGBoost. First column compares treated units to individuals that received Pandemrix (control units). 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.

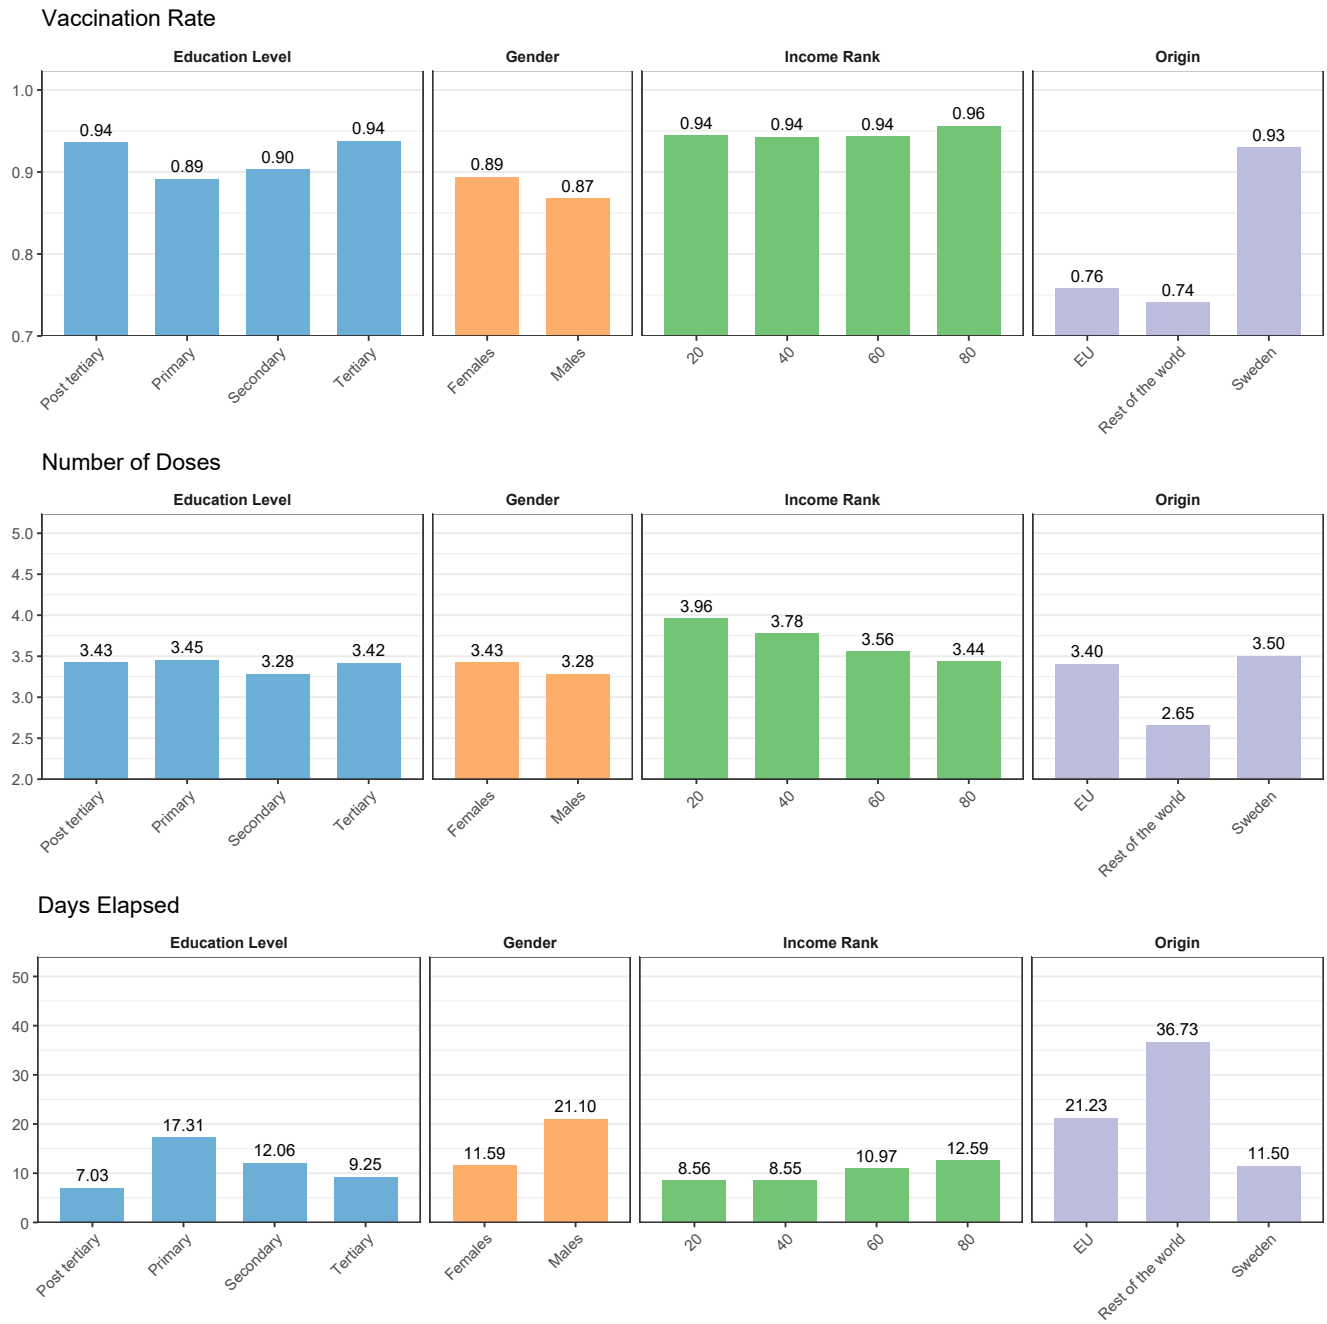

Figure B1: COVID-19 immunization outcomes for different sub-populations.

Notes: Each row reports COVID-19 immunization outcomes for different sub-populations. **First column:** by education level (individuals aged 40–50 in 2022). **Second column:** by gender (individuals aged 40–50 in 2022). **Third column:** by income percentile (individuals aged 40–50 in 2022). **Fourth column:** by origin (Swedish vs. foreign background; foreign defined as being born abroad or having two parents born abroad; individuals born 2004 or earlier). Outcomes include vaccination rate (top row), average number of doses (middle row), and average days elapsed since first eligibility (bottom row).

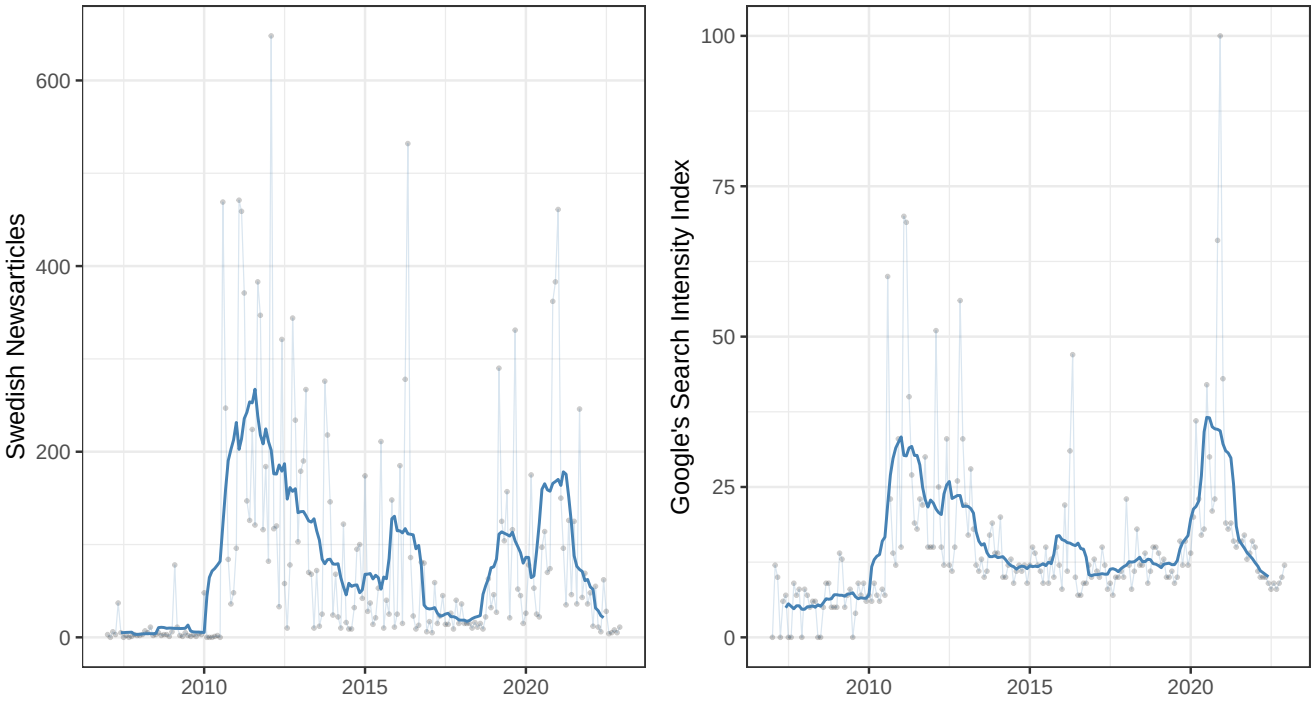

Figure B2: Exposure and Attention to the Narcolepsy Scandal

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

Figure B3: Narcolepsy Drug Prescriptions Across Time

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 psychoanaleptics: 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.

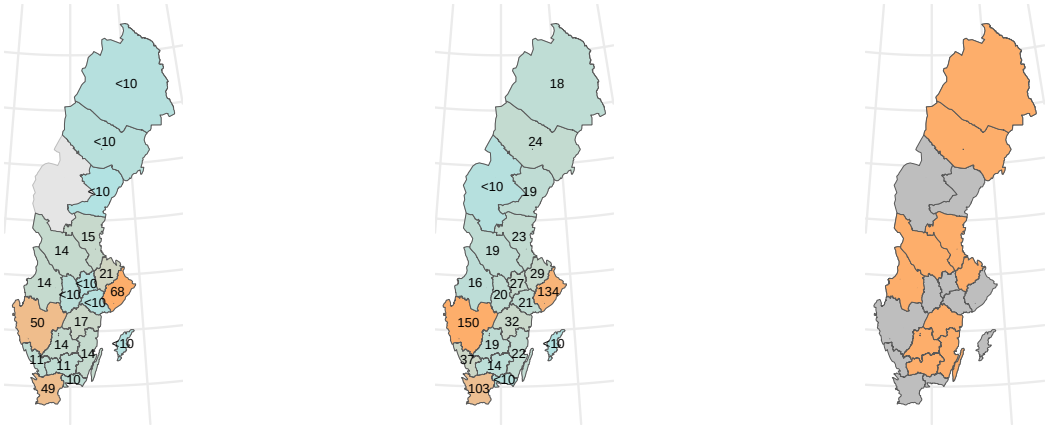

(a) Post–Swine Flu Cases **(b)** Pre–Swine Flu Cases **(c)** Pandemrix Vaccinations

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

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

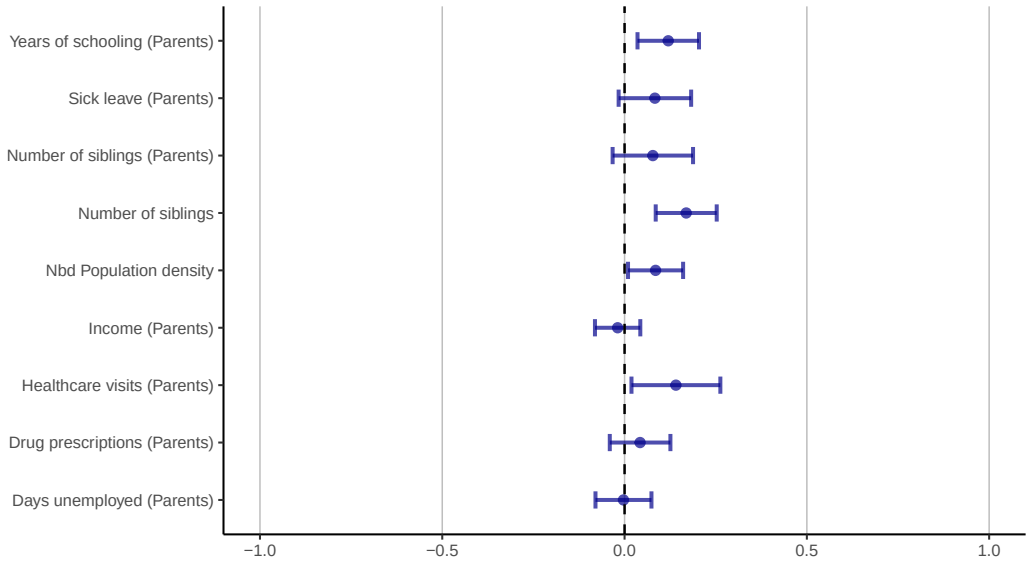

Figure B5: Balance in Pre-treatment Covariates, Pre-Swine Flu Sample

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 B6: Balance – Alternative Matching Methods

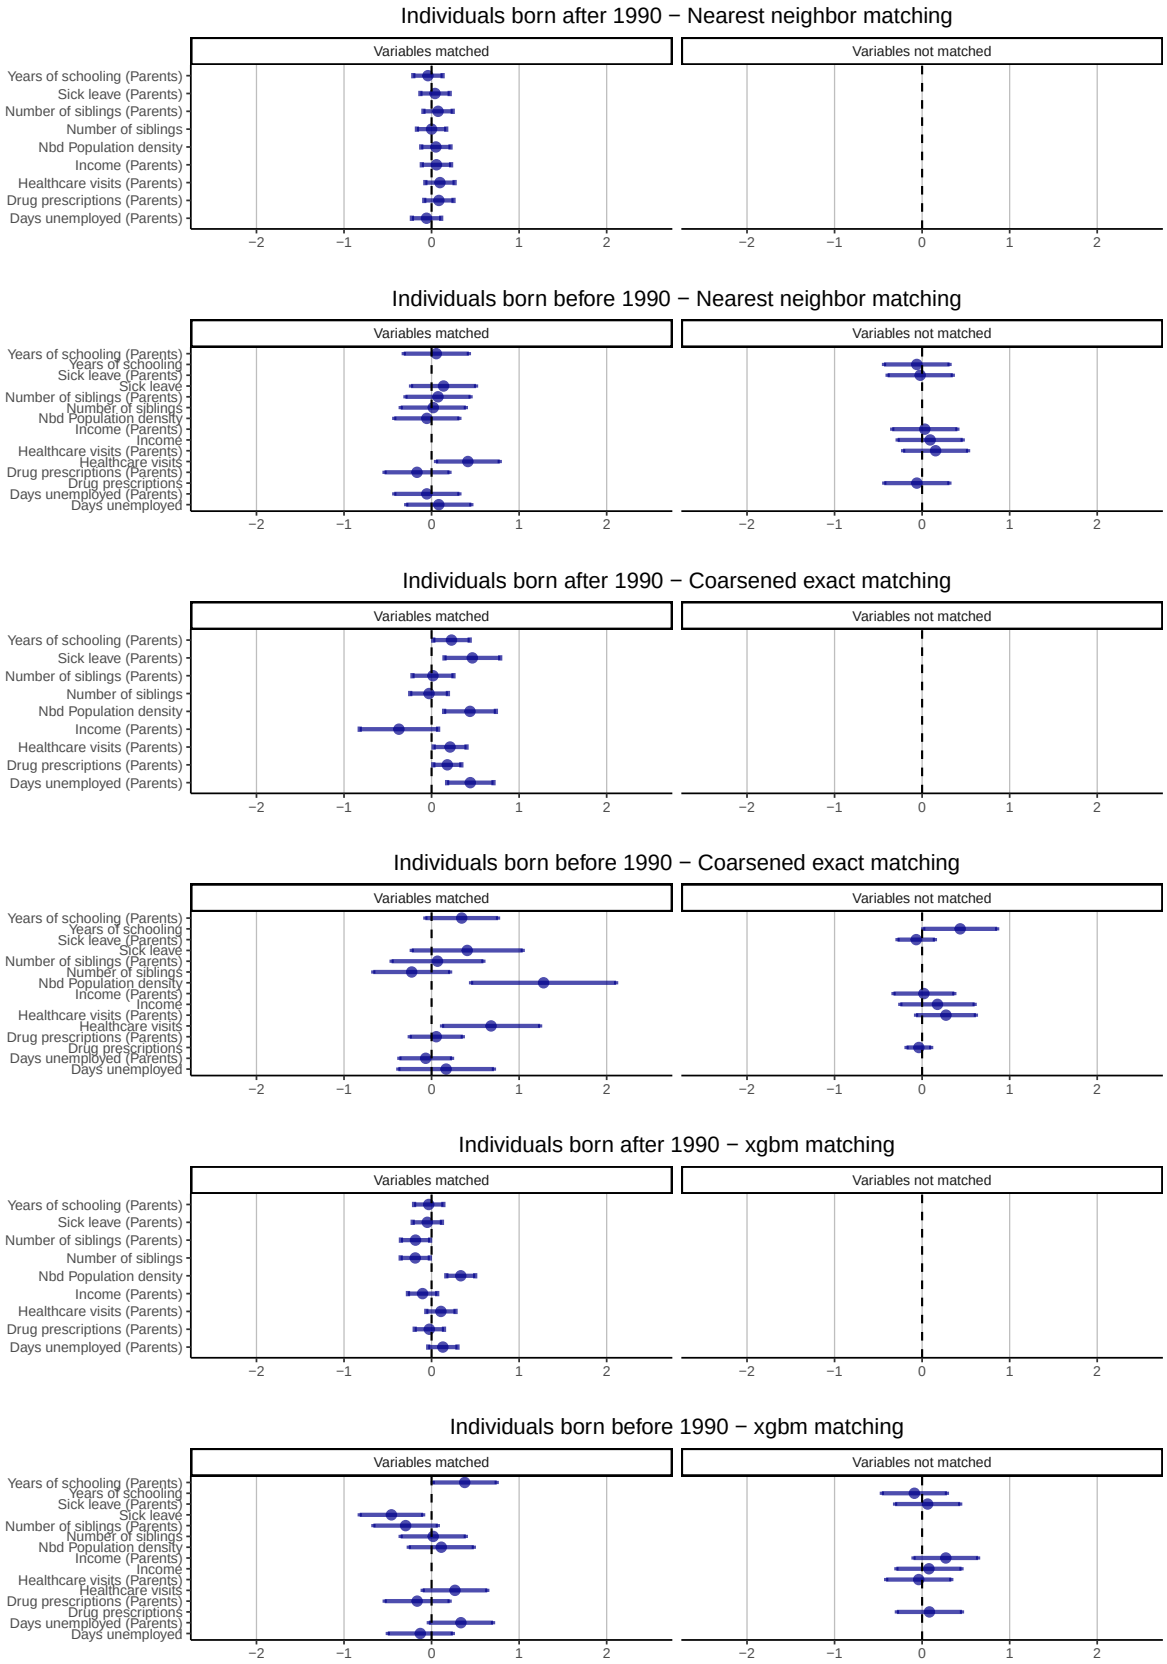

Notes: This figure displays differences in pre-determined characteristics between individuals that develop narcolepsy and control individuals for the main sample. Each estimate comes from a univariate regression of the standardized covariate on the treatment status. **Left subfigures:** Differences between treated and control units in the matched sample for matched variables. **Right panels:** Differences between treated and control units in the matched sample for variables not chosen for matching. Note that the variables used for matching is independent of the matching method.

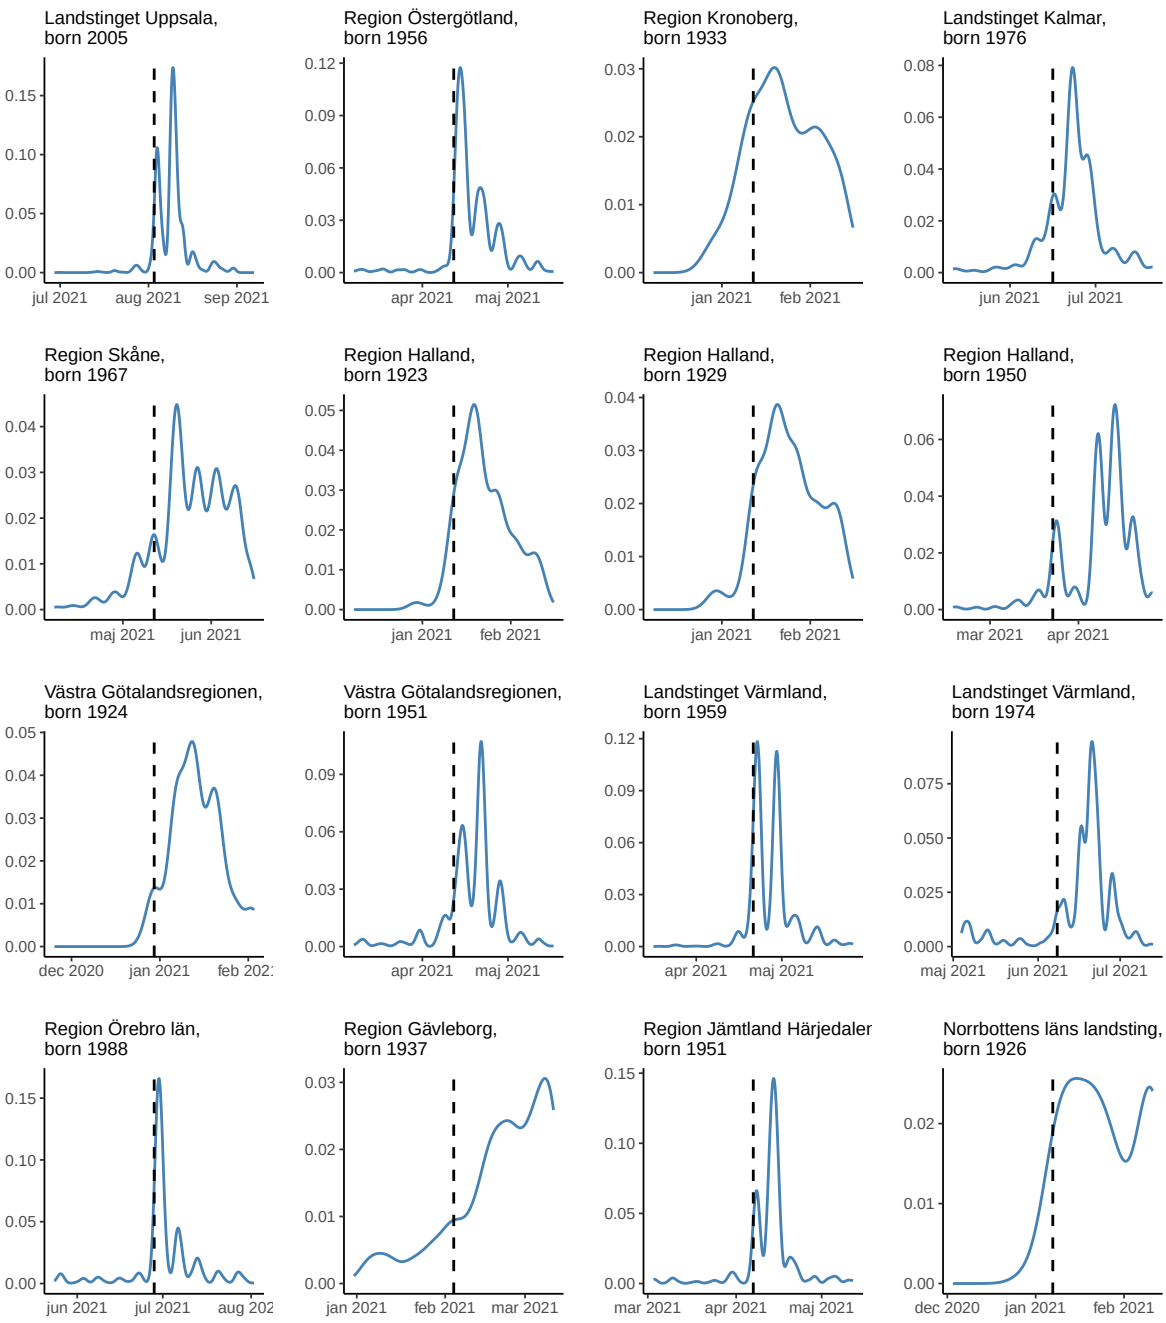

Figure B7: Definition of First Date of Availability

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

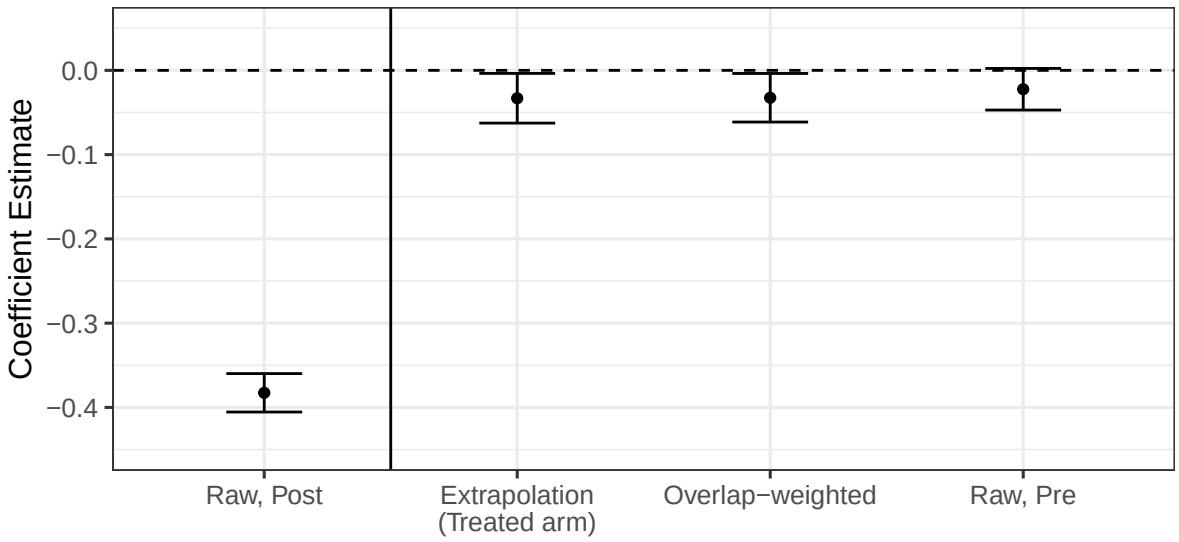

Figure B8: Counterfactual Narcolepsy Effect

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–weighted** extrapolate/predicts vaccine uptake to the to the age distribution of post swine flu individuals from the pre swine flu-sample. **Raw, Pre** is the unadjusted treated–control difference in the Pre-Swine Flu sample.

Figure B9: Main results – Partners.

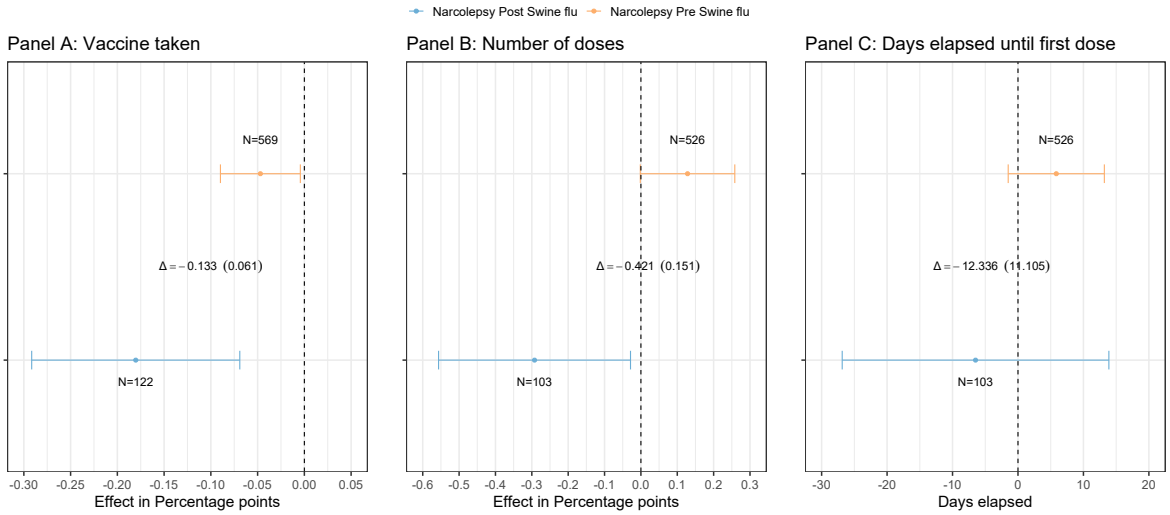

Notes: This figure displays coefficient corresponding to τ_{post} and τ_{pre} in eq. (4), similar to fig. 6 but for partners of the focal individuals. **Panel A:** Displays results for whether or not an individual has taken at least one COVID-19 dose. **Panel B:** Displays results for the number of COVID-19 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, i.e. by the treated or control individual that a network member is related to.

Table B2: Immunization outcomes and Mild Pandemrix Side Effects.

	Vaccine Taken	Number of Doses	Days Elapsed
Treated	-0.036	-0.061	-4.593
	(0.008)	(0.025)	(1.847)
Num.Obs.	3907	3648	3648
R2	0.055	0.273	0.045

Notes: This table displays estimates from regressing each outcome variable on a binary indicator equal to 1 if an individual reported a mild side effect from Pandemrix. Each treated individual is matched 1:1 to a control individual who also received Pandemrix but did not report any side effects. Matching is based on propensity scores estimated using pre-treatment health and socioeconomic characteristics. Matches are restricted to individuals of the same birth year and gender.

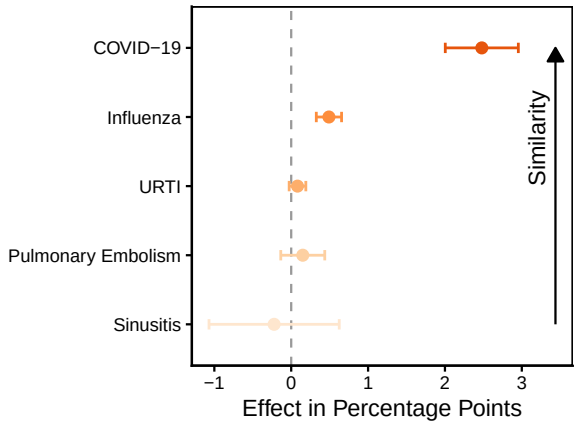

(a) Vaccine Taken

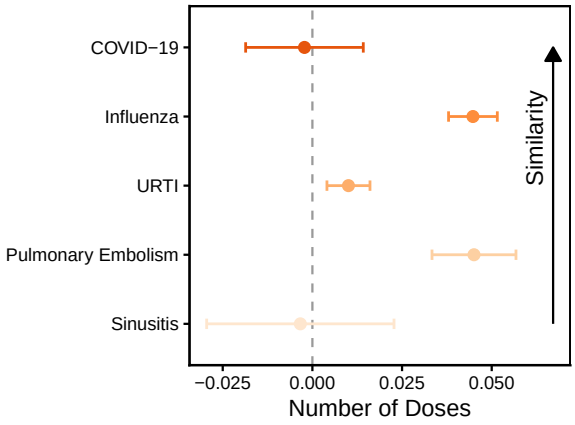

(b) Number of Doses

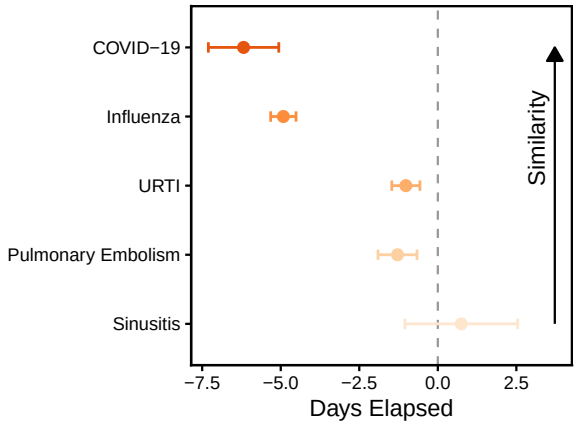

(c) Days until First Dose

Figure B10: Benefits from Vaccines.

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.

Table B3: Adverse Events and Immunization Outcomes – Age Reweighted

	General Side Effects			Vaccine Side Effects		
	Vaccine Taken	Number of Doses	Days Elapsed	Vaccine Taken	Number of Doses	Days Elapsed
Diseased Individuals						
Side Effect	-0.019	-0.025	-1.90	-0.097	-0.062	11.2
	(0.003)	(0.013)	(0.63)	(0.013)	(0.059)	(2.95)
N. Treated	15 294	13 974	13 974	849	686	686
Family Members						
Side Effect	-0.001	0.048	0.74	-0.042	-0.15	7.83
	(0.002)	(0.009)	(0.38)	(0.008)	(0.037)	(1.54)
N. Treated	34 018	31 445	31 445	2 048	1 771	1 771
Children						
Side Effect	0.005	0.015	-0.77	-0.029	0.072	3.25
	(0.003)	(0.008)	(0.52)	(0.018)	(0.044)	(2.89)
N. Treated	28 087	22 124	22 124	887	662	662
Partner						
Side Effect	-0.003	-0.015	-0.78	-0.036	-0.12	5.03
	(0.003)	(0.018)	(0.72)	(0.014)	(0.086)	(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 a side effect between 2015 and 2020. Columns 1–3 use an indicator for reporting any side effect that was not vaccine-related between 2015 and 2020 as the regressor ("General side effects"). Columns 4–6 instead use an indicator for reporting a vaccine side effect ("Vaccine side effects"). 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.

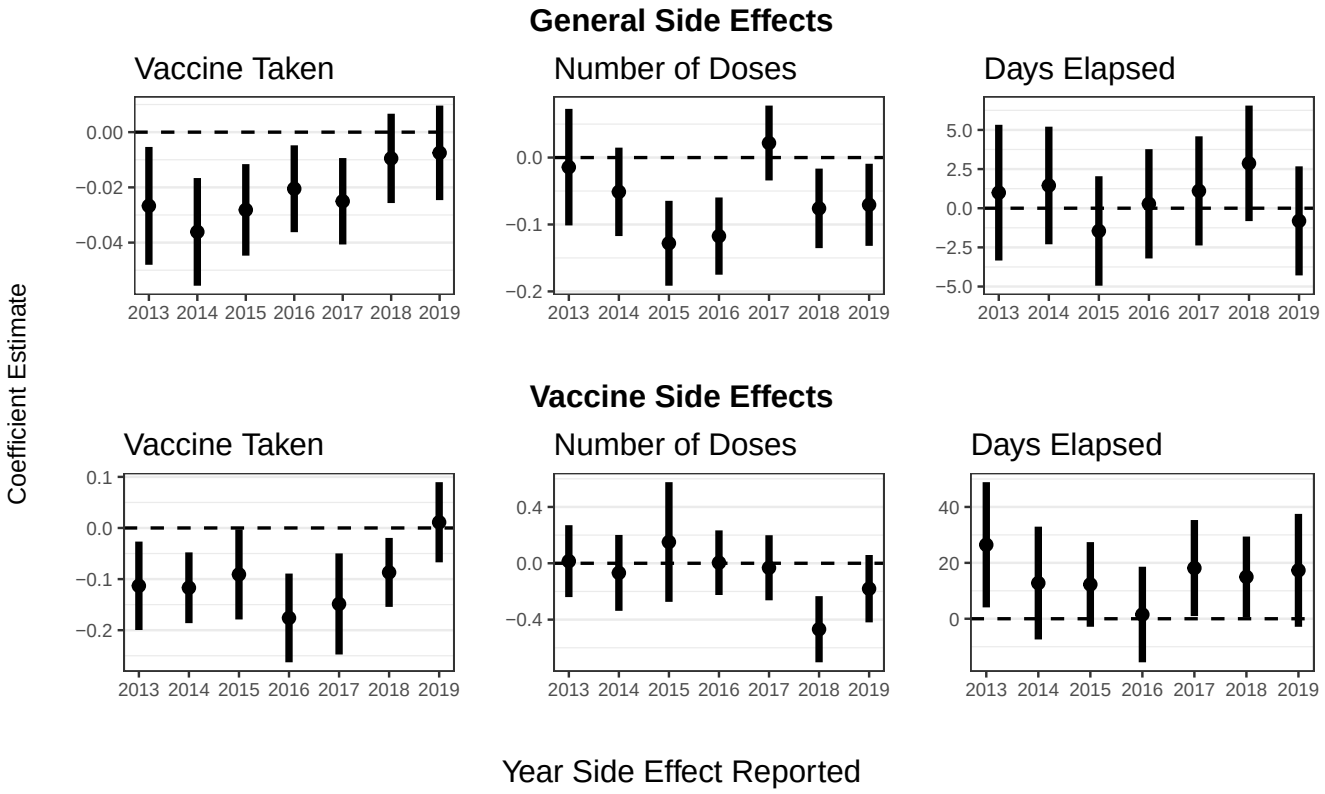

Figure B11: Adverse Events and Immunization Outcomes – Heterogeneity by Time.

Notes: This figure displays results from regressing the COVID-19 outcome variables on general side effects and vaccine side effects split up by year of reporting the side effect. Observations are reweighted to match a common age and drug distribution across years.

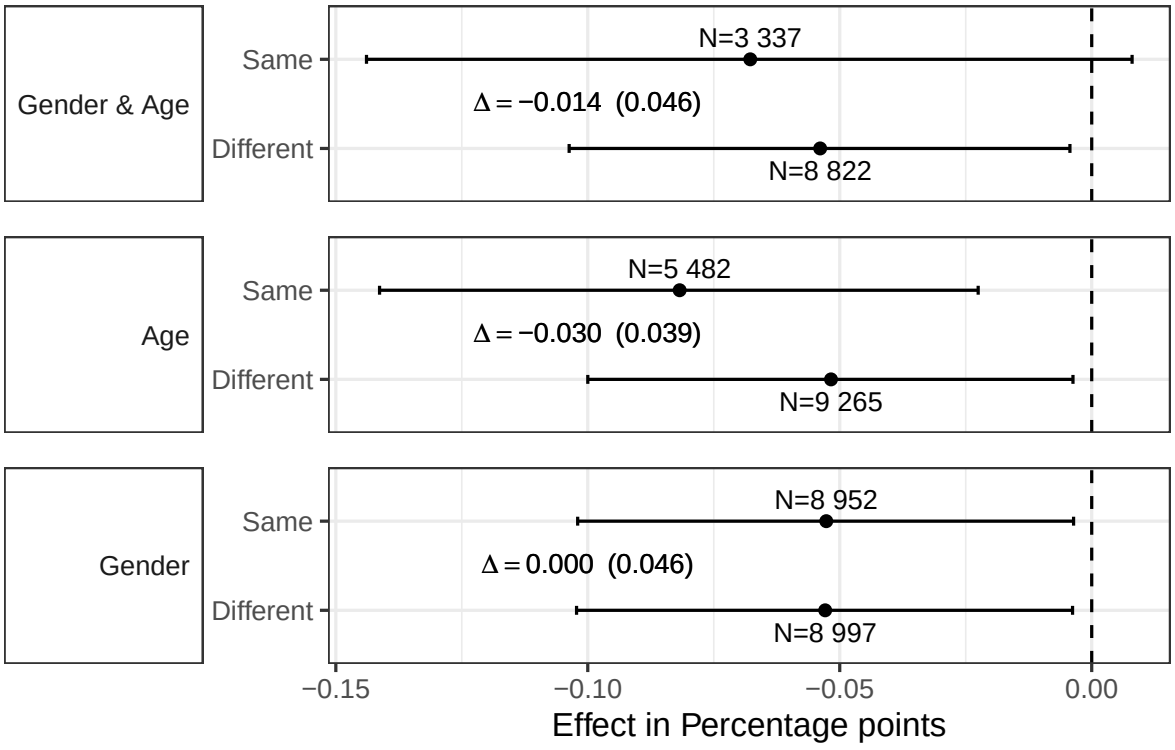

Figure B12: Heterogeneity—Similarity-based learning

Notes: This figure displays the estimated coefficients on τ_{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.

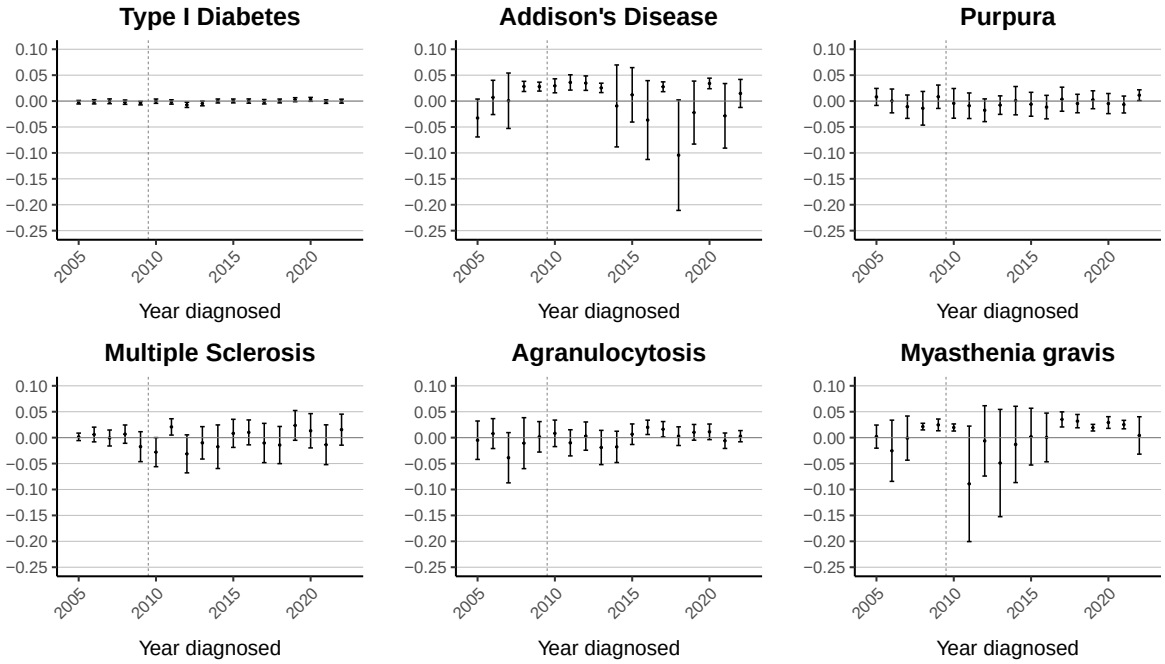

Figure B13: Other Severe Diseases—Heterogeneity Across Time

Notes: This figure displays the estimated coefficients on τ_{post} for the outcome variable *Vaccine Taken* for other diseases that people may have attributed to Pandemrix. The control group consist of individuals who received Pandemrix but did not developed narcolepsy. The treatment group consist of individuals who were first diagnosed with each of the respective diagnoses in year t. Each specification includes birth year fixed effects.

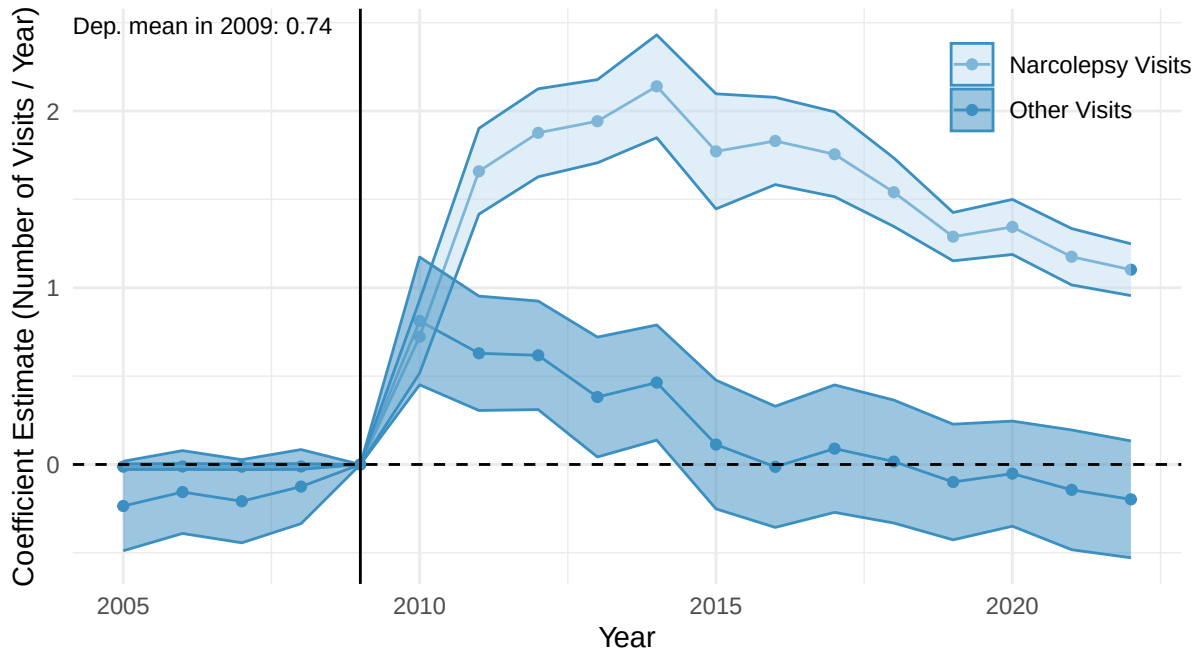

Figure B14: DiD-Estimates of Healthcare Visits After Narcolepsy Diagnosis

Notes: This figure displays difference-in-differences estimates from regressions where the outcome variable is the number of healthcare visits per year. The sample consists of treated individuals matched on birth year to others who also received Pandemrix. Each treated individual is matched to three control units. 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 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 exanthems 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

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

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

Appendix D. Regional Pandemrix data

Table B4 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 side effects are self reported, the remaining reports are made by doctors and nurses. In our preferred specification we include both individuals reported side effects themselves and individuals that had a doctor or a nurse report the side effect for them.

Table D1: Reported mild side effects.

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 side effects 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.
Norrbotten	131,195	0.53	
Uppsala	181,461	0.54	
Värmland	116,931	0.43	
Västerbotten	3,525	0.01	Coverts vaccinations after 2010-10-01 only.
Östergötland	274,405	0.64	

Notes: Vaccination rates are calculated as the share of individuals residing in each healthcare region in 2009.

Appendix E. Robustness of Main results

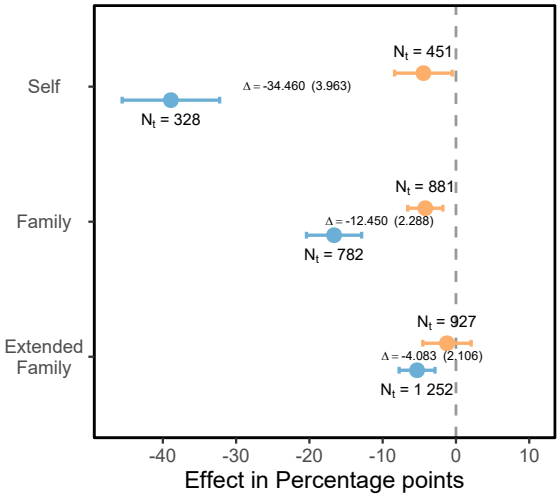

(a) Vaccine Taken

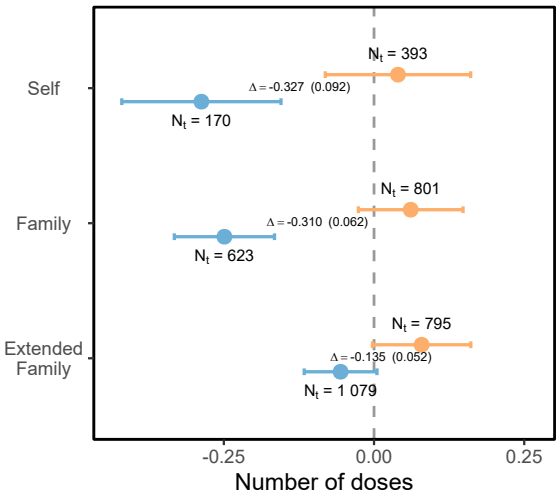

(b) Number of Doses

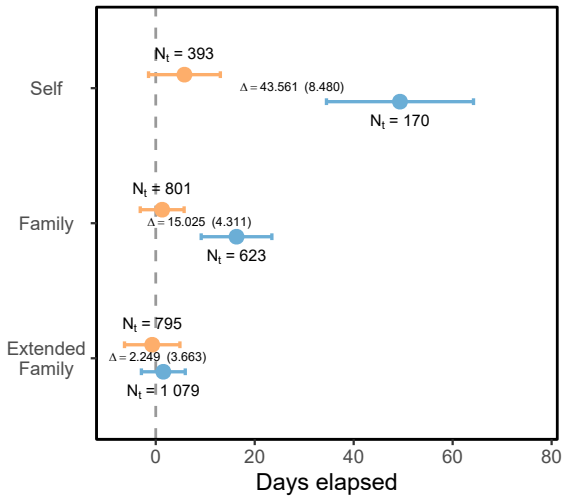

(c) Days until First Dose

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

Notes: This figure displays coefficients corresponding to τ_{post} and τ_{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, i.e. by the treated or control individual that a network member is related to.

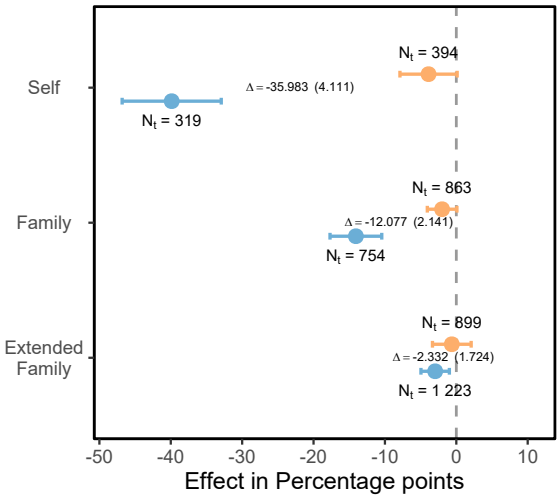

(a) Vaccine Taken

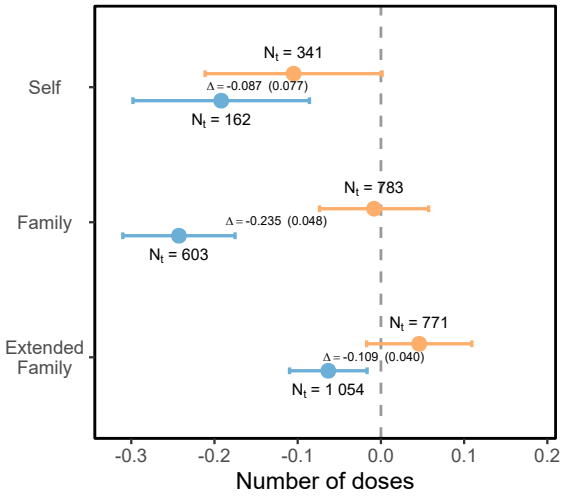

(b) Number of Doses

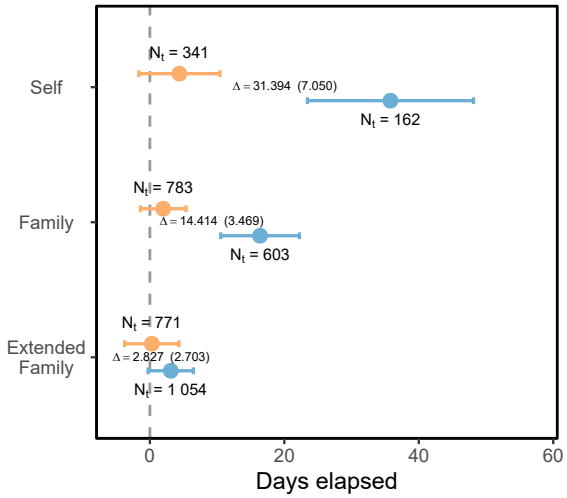

(c) Days until First Dose

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

Notes: This figure displays coefficients corresponding to τ_{post} and τ_{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, i.e. by the treated or control individual that a network member is related to.

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

Model hyper-parameters	
Maximum tree depth (`max_depth`)	4
Learning rate (η)	0.10
Minimum child weight (`min_child_weight`)	1
Subsample fraction used for training (`subsample`)	0.8
Number of boosting rounds (`n_rounds`)	100
Model setup	
Cross-validation folds	5
Number of treated units	1 013
Features	
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 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 (batch size 256, 5 epochs). The resulting sequence embedding is concatenated with imputed and encoded socioeconomic covariates and fed into an XGBoost classifier trained with logistic loss. Final performance is evaluated on the 10% hold-out test set.

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

Sequence encoder (RNN)	
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)
Gradient boosting classifier (XGBoost)	
Objective	`binary:logistic` (logistic loss)
Boosting rounds (`n_estimators`)	4000 (with early stopping)
Learning rate (η)	0.02
Max depth	5
Subsample / colsample_bytree	0.8 / 0.8
Regularization	$\lambda = 1.0$, $\alpha = 0.0$
Class imbalance	`scale_pos_weight` = $n_{\text{neg}}/n_{\text{pos}}$
Early stopping	200 rounds on a 10% validation slice of training data
Model setup	
Train / test split	90% / 10% (stratified)
Hybrid features	Concatenate sequence embedding (32-D) with tabular covariates
Feature sets	
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 side effects (sequence input)

Appendix G. Mathematical Derivations

MAR among treated

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

yields that

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

similarly, MAR among untreated implies that

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

MAR(1) and source irrelevance directly give us that

$$E\big[Y(1) \mid D = 1, X] = E[Y(1) \mid W = 1, D = 1, X] = E[Y(1) \mid W = 1, X]$$

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

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

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

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

We maintain the assumption that $\Pr\big(W = 1 \mid X\big)$ is small relative to the number of non-reporters such that the term vanishes. Together with mar(2) it directly follows that

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

Appendix H. Classifying Adverse Events into Type-Learning and 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 the following 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 a side effect teaches me about my likelihood of experiencing a side effects. 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. Finally, we classify each combination as idiosyncratic if its score falls below 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

Idiosyncratic	**Type-learning**
Psychostimulants — Fluctuating mood symptoms	**Penicillins** — Allergic conditions NEC
Aspirin — Nasal disorders (epistaxis)	**Enalapril** — Angioedemas
Levothyroxine — Asthenic conditions	**Sirolimus** — 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.