Post Hoc Ergo Propter Hoc?

Side Effect Misattribution and Vaccine Hesitancy

Svenja Miltner* Jonatan Riberth† Anton Arbman Hansing‡

This version: April 8, 2026

Abstract

Side effect concerns are the most frequently reported reason for vaccine hesitancy, yet we lack well-identified evidence on how these risk perceptions form. Using Swedish register data, we study how experiencing illness shortly after vaccination affects the decision to continue vaccinating. We exploit an episode during the COVID-19 vaccination campaign when links between the AstraZeneca vaccine and rare blood clots triggered extensive media coverage and temporary suspensions. Individuals who experience a common blood clot shortly after a dose—clinically unrelated to the vaccine-induced syndrome—drop out of vaccination at almost twice the rate of carefully matched controls (+4pp). Consistent with misattribution, the effect is strongest for diagnoses occurring soon after the dose, and greatly amplified (+17pp) when clinicians report the event as a suspected side effect. Other acute conditions of similar severity produce no comparable response. Strikingly, effects are similarly strong for mRNA vaccine brands, which operate through a different mechanism and were never scientifically linked to vaccine-induced blood clots. Two groups respond most strongly: younger individuals (who face lower COVID-19 risk) and those who delayed their first dose (indicating baseline skepticism). These findings point to an underappreciated tradeoff: while transparency about vaccine safety is thought to build public trust, such signals may also amplify perceived risks beyond what the evidence warrants.

Keywords: COVID-19, Side Effects, Vaccine Hesitancy

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

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

‡IIES, Stockholm University. E-mail: Anton.Hansing@iies.su.se.

The authors are grateful to Arash Nekoei, Torsten Persson, Anna-Mia Ekström, Tove Fall, Olof Östergren, Anna Dreber-Almenberg, Natalia Danzer, Peter Mohr, participants at the 2025 SWECOV workshop and the fall PhD workshop at Freie Universität Berlin. 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

Vaccine hesitancy—the delay or refusal of available and effective vaccines—remains a critical challenge in containing communicable diseases, and is often driven by fears of harmful side effects (Jones, 2020). Personal experiences shape such beliefs and behaviors far beyond their informational content. Whether an experience shifts behavior depends on how individuals understand its cause. For example, when symptoms arise after vaccination, individuals may infer a link between the vaccine and their health state even when the two are unrelated. Such misattribution can be etched into personal memory and distort perceived risks, discouraging future vaccination. This paper studies when and for whom misattribution occurs, and how it ultimately affects vaccination behavior.

How individuals interpret health events depends on the informational environment in which they occur. In vaccination, a key part of that environment is communication about suspected side effects. When these are flagged by expert agencies, media, or clinicians, individuals may overestimate the risks of vaccination. As information spreads increasingly quickly and widely, understanding how such communication affects causal attribution and risk perceptions is important for pandemic management.

We examine how personal exposure to health events after vaccination — and the informational environment surrounding them — affects subsequent vaccine uptake. Our setting is the AstraZeneca COVID-19 vaccine, which in early 2021 was linked to a rare blood clot, triggering widespread media coverage and temporary suspensions in several countries. Using Swedish register data on individual vaccine decisions and diagnoses, we show that experiencing an unrelated blood clot after vaccination substantially increases dropout from the COVID-19 vaccination schedule. The effects are larger when the blood clot occurs soon after vaccination and when a clinician reports it as a suspected side effect. Other acute conditions, whether plausibly vaccine-induced or not, do not generate similar responses.

Sweden commenced vaccination against COVID-19 in late 2020 and predominantly administered two mRNA vaccines, by Pfizer and Moderna, and one adenoviral vaccine, by AstraZeneca. In March 2021, reports from several countries raised the suspicion of a possible link between the AstraZeneca vaccine and Thrombosis with Thrombocytopenia Syndrome (TTS), an extremely rare but severe blood clotting condition. Over the course of a few days, several countries suspended the AstraZeneca vaccine pending safety reviews. Though the early estimates of TTS incidence ranged between 10 to 20 cases per million administered doses (MHRA, 2023), the extensive media coverage could have contributed to public skepticism about vaccine safety related to unrelated blood clots—conditions affecting over 0.5% of the Swedish population annually. By the time it was suspended in Sweden on March 16, the AstraZeneca vaccine had been in place for three months with a total of 230,000 doses administered. The Swedish suspension was eventually lifted and vaccination resumed on March 25 for individuals above the age of 65, until the vaccine was finally phased out by July 2021.

Our empirical strategy compares individuals with a similar COVID-19 vaccine history, where some experienced a blood clot plausibly unrelated to the vaccine soon after vaccination while others did not. The credibility of this design stems from the close similarity achieved by pairing each treated individual with a control who received the same vaccine brand(s) in the same month(s), and who is otherwise similar in

demographic, socioeconomic, and health characteristics that may shape both blood clot risk and subsequent vaccine uptake.

Our findings suggest that individuals generalize signals about specific but salient adverse events to other similar diagnoses and drugs. In particular, we find that experiencing a blood clot, unrelated to the COVID-19 vaccine, within six weeks of vaccination reduces uptake of the next dose by 4 percentage points—a 70% increase of the dropout risk from the vaccination program relative to a baseline uptake rate of 94.5% in the control group. This effect is plausibly driven by by external cues through intense public attention that cause individuals to attribute unrelated symptoms to the vaccine they received, despite the distinct pathology of vaccine-induced blood clots. We find moderate spillover effects on close family members. Combining the response of the diseased individual with these spillovers, roughly half of the overall increase in vaccine hesitancy comes from the individual who experienced the blood clot, while the other half comes from close family members who subsequently discontinue vaccination. We find only weak evidence to support that a general misattribution mechanism is at play, where unrelated acute events that received no media coverage, other than blood clot, are also routinely attributed to recent vaccination. In particular, for other acute conditions that could plausibly have been attributed to the vaccine, we find only very modest effects on dropout rates.

Two features of the data indicate that the dropout effect is proportional to the strength of the informational signal. First, we find that the effect size is greatest for blood clots suffered soon after the vaccination, when the causal link to the dose taken is most plausible. If the dropout effect were instead due to medical incapacitation, we would expect the opposite pattern, as individuals diagnosed earlier have more time to recover before their next dose is due. Second, when a clinician reports the blood clot event as a side effect of the vaccination—a choice that is often communicated to the patient—the vaccine uptake is reduced by a further 14 percentage points. The total effect on these individuals is a close to five-fold increase in dropout rates compared to matched controls. Though we cannot rule out some selection into reporting, this suggests that clinician communication—much like media coverage—cues individuals to believe that the side effect was real.

Notably, we find that the effect is of similar magnitude across vaccine brands, despite no suspected link for vaccines other than AstraZeneca. We infer that, even if prevailing information suggests a specific vaccine-adverse event relationship, individuals do not discriminate across vaccine brands in their risk perception of adverse events when deciding to complete their vaccination schedule.

We proxy perceived vaccination benefits using patients' age and how early they were vaccinated relative to when the vaccine became available, and find smaller effects among individuals with higher expected benefits from vaccinating. This pattern is consistent with vaccination decisions reflecting a trade-off between perceived benefits and perceived costs in the form of side effects. It is less consistent with alternative explanations in which uptake falls mechanically due to incapacitation, or in which the blood clot episode triggers a non-compensatory reaction such as generalized distrust under which high-benefit individuals would be no less likely to drop out.

Finally, we find no heterogeneity with respect to three proxies for health literacy: high cognitive ability, having a university degree and having a medical practitioner in the family. Individuals with higher health literacy are equally likely to attribute blood clots to the vaccine they received.

Our paper speaks to both behavioral economics and public health. First, we relate to a literature on how individuals use past experiences in decision making (Ashraf et al., 2024; Malmendier et al., 2021), including recent work that studies how such experiences shape vaccination decisions (Bordalo et al., 2022). We contribute by focusing on high-stakes choices during a pandemic and by documenting how salient adverse events, accompanied by informational cues, spill over across health domains.

Second, we contribute to a literature on how individuals infer causal relationships from the temporal proximity of events, even when evidence of an underlying link is weak or absent. While well established in cognitive psychology—often described as an "illusion of causality"—this phenomenon and its consequences have received limited attention in economics (see, e.g., Malmendier and Tate (2005); Espín-Sánchez et al. (2023)). Our focus is on the extent of such misattribution in the healthcare domain.[1]

We add to a small literature studying how reports linking the AstraZeneca COVID-19 vaccine to blood clots shaped vaccine attitudes during the pandemic. Agosti et al. (2022) and Deiana et al. (2022) analyze vaccine attitudes around the suspension of AstraZeneca. Using aggregate-level data across Europe, they provide evidence that the AstraZeneca blood-clot episode, and particularly the subsequent retraction, fueled skepticism extending beyond the AstraZeneca vaccine itself. While these studies document cross-vaccine spillovers, our rich individual-level data allows us to how risk perceptions are shaped at the individual level by exploiting personal exposure to blood clots. Furthermore, we observe actual vaccination decisions rather than stated preferences or intensions as measured through surveys.

The paper proceeds as follows. Section 2 provides a background on vaccine-induced blood clots and the AstraZeneca suspension. Section 3 develops a framework predicting which health signals drive vaccine dropout and who responds most. Section 4 introduces the register data and defines key variables used to test these predictions, and Section 5 details the matching strategy. In Section 6 we present our findings, while Section 7 discusses implications for risk communication.

2 Vaccine-Induced Blood Clots

Sweden rolled out its first COVID-19 vaccines in December 2020 and, by February 2021, had deployed vaccines from all three major suppliers in Europe: BioNTech/Pfizer, Moderna, and Oxford–AstraZeneca. Soon after market access, the EMA received an unexpected number of reports of severe side effects following vaccination with AstraZeneca. Typically, patients would show an unusually low count of blood platelets, internal bleeding, and blood clots forming at unusual sites in the body, such as the brain's dural venous sinuses. By March 16, three Northern European countries suspended AstraZeneca's COVID-19 vaccine Vaxzevria alongside Sweden within a week's time.[2] Pending a formal investigation of the events, the EMA announced that the overall benefits of vaccination still outweighed risks of suffering a rare blood clot. As

[1] The illusion of causality is related, but different from nocebo responses in the medical literature, where expectations about adverse effects lead individuals to experience (psychosomatic) symptoms even in the absence of an active ingredient. A substantial share of participants in the placebo arms of COVID-19 vaccine trials report adverse events following vaccination (Haas et al., 2022). Pertinently, Asan et al. (2024) show that patient reports of severe headache increased substantially after media coverage of rare vaccine-induced blood clot cases and associated warning symptoms in Germany.

[2] The vaccine held market access authorization for the entire European Union and several other countries. Globally, in March 2021, more than 20 countries halted the administration of the vaccine within a matter of days.

a result, the Public Health Agency of Sweden resumed vaccination with AstraZeneca for people aged 65 and over on March 25. The EMA finally confirmed a causal link between the Vaxzevria vaccine and the extremely rare and potentially fatal condition known as vaccine-induced thrombosis with thrombocytopenia (VITT). Thereafter, demand for AstraZeneca's vaccine in Sweden collapsed, and by May 2021 it was phased out as the rollout continued almost entirely with mRNA vaccines.

While the pathological mechanism of VITT is still under investigation, current studies suggest that VITT is caused by an autoimmune response to components specific to adenoviral vaccines. This is in line with the observation that the adverse events are predominantly observed in patients vaccinated with J&J and AstraZeneca, but not with common mRNA vaccines. Further evidence suggests that any increased thrombotic risk after AstraZeneca is concentrated in rare TTS cases, rather than reflecting a broad increase in thrombotic events.

Note that the unusual clinical picture of VITT, including diagnosis criteria of platelet formation and widespread locations of clotting, allowed clinicians to quickly rule out VITT in the vast majority of post-vaccination blood clots once diagnostic criteria for VITT were established. Although differences in data acquisition between countries lead to varying estimates of incidence, the European Medical Agency suggested that on aggregate VITT occurred in about 1 in 150 000 individuals vaccinated with AstraZeneca (including UK cases) (Klok et al., 2022; Franchini et al., 2021). To benchmark this against other common types of blood clots: Pulmonary embolism, another severe clotting event, has an incidence rate of 167:100 000 among the Swedish population (Johansson et al., 2014).

Evidence on who was at risk of developing VITT is limited, but perceived risk may nonetheless have varied across groups. UK safety monitoring suggested that the event was twice as common among individuals aged under 50 (about 1 in 50,000) as compared to those above the age of 50 (1 in 100,000), and early case reports appeared to involve more women (MHRA, 2023). Further investigations revealed, that these patterns may partly have reflected differential rollout and reporting rather than personal risk factors. Ultimately, studies have not identified individual characteristics that predict VITT (Klok et al., 2022).

Experiencing a blood clot does not constitute a medical contraindication to COVID-19 vaccination. For example, public health authorities did not advise individuals with a history of common blood clots, such as deep-vein thrombosis following surgery or during pregnancy, to avoid vaccination. However, individuals who developed a blood clot soon after vaccination with AstraZeneca were recommended to complete their vaccination schedule with an mRNA vaccine. Furthermore, the general scientific consensus is that any vaccination against COVID-19 is still highly beneficial even given the risk of general thrombotic events. In particular, studies have shown that the risk of suffering a blood clot is significantly higher following infection than following vaccination (Zhao et al., 2024; Katsoularis et al., 2022).

Our empirical strategy builds on reports of vaccine-related blood clots serving as vaccine-safety cues under broad public attention. The mechanism is an interaction: public reporting is common to all, but personal exposure makes clot-related news more salient and more likely to affect behavior. A necessary precondition is that the controversy received broad public attention. Previous studies have documented a spike in public interest during the time around the AstraZeneca suspension, using Google searches as a proxy for awareness (Deiana et al., 2022; Agosti et al., 2022), suggesting that the adverse events were part of

a public debate. In Sweden, we observe a similar pattern: Google searches related to blood clots spiked in March, when national health authorities announced the suspension of AstraZeneca (see Figure B1). In Figure B2 we provide further evidence on the attention that the controversy received and on whether this attention was not specific to the AstraZeneca vaccine. In particular, we plot the number of side effects for the three major vaccines against the number of administered doses. On the day of suspension and shortly after, reported side effects increased substantially for AstraZeneca and also Pfizer, which is commonly perceived as a safe option and has not come under scrutiny in the public debate.[3]

3 Conceptual Framework

Motivation This section presents a simple model of how vaccination choices are impacted by the experience of perceived side effects. Individuals begin by choosing whether to get an initial vaccine dose, weighing benefits against the perceived probability and severity of side effects. Following the initial dose, they may experience a health event. Drawing on the type of illness, the time from vaccination to onset, and possible links drawn between vaccination and the symptoms by media and doctors, the individual estimates a probability that their ailment was caused by the vaccine. They then update their risk assessment of the vaccine before deciding whether to take the next dose.

The model, while simple, delivers predictions about which kinds of ailments and information treatments cause individuals to drop out of the vaccination program, as well as which individuals are most responsive. We bring these predictions to the data.

Learning about vaccine risk Individuals are uncertain about the true probability of serious vaccine side effects $\theta \in [0, 1]$. Before vaccination, their belief over θ is captured by a shared prior $p(\theta)$. After receiving the first dose, individuals may or may not experience a vaccine-induced illness—$z_i \in \{0, 1\}$—where $z_i = 1$ with probability θ.

Importantly, z_i is unobserved by the individual. Instead, she observes a health signal y_i. The signal consists of (i) symptoms (or lack thereof), (ii) time from vaccination to onset, and (iii) whether the condition has been linked to the vaccine by media coverage or the treating clinicians. The distribution of y_i depends on whether the illness was truly vaccine-induced ($z_i = 1$) or not ($z_i = 0$). We denote the respective conditional likelihood functions by g_1 and g_0. g_0 can be thought of as the distribution of health outcomes absent vaccination, whereas g_1 represents the individual's view of how vaccine side effects tend to present.[4]

The individual uses the health signal y_i to form a posterior distribution over θ:

$$p(\theta|y_i) \propto p(\theta) \cdot (\theta \cdot g_1(y_i) + (1 - \theta) \cdot g_0(y_i)) \tag{1}$$

which is then used to decide whether to take the next dose.

[3] Note that although not as prominently discussed in public, the mRNA vaccines Pfizer and Moderna have been linked to rare cases of myocarditis and pericarditis.

[4] Concretely, let $y_i \in \{\emptyset\} \cup \mathcal{Y}$ where $\mathcal{Y}$ is the set of health events and $\emptyset$ denotes remaining healthy. Then $g_0(\emptyset)$ will be relatively high whereas $g_1(\emptyset) \approx 0$. The degree to which an event y is attributed to the vaccine is captured by the likelihood ratio $g_1(y)/g_0(y)$. The ratio is high for symptoms that occur soon after vaccination or matching the individual's prior about how a side effect present. It is low for illnesses with delayed onset or no plausible causal link to the vaccine— $g_1(\text{broken arm}) \approx 0$.

Vaccination choice Individuals base vaccination choices on beliefs about side-effect risks, weighed against benefits. We denote the perceived benefits of vaccination by $B(x_i)$, where x_i is a vector of observable traits. These benefits include both health and non-health considerations.

The costs of vaccination consist of the perceived risk of side effects, $E[\theta]$, multiplied by a shared severity S. For the first dose, the expectation over θ is computed with regards to the prior distribution $p(\theta)$, whereas for the second, it is computed from the posterior $p(\theta|y_i)$. Finally, we allow for an idiosyncratic shock ε_i to account for attitudes toward vaccination not captured elsewhere. The perceived net benefit of vaccination is

$$u_i = B(x_i) - E[\theta] \cdot S + \varepsilon_i \tag{2}$$

Individuals choose not to vaccinate if $u_i \leq 0$. If $u_i > 0$, individual i gets vaccinated t_i days after the vaccine is made available to them, where $t_i|u_i \sim G(\cdot|u_i)$ such that higher u_i implies earlier vaccination in distribution[5].

Model predictions The model delivers five predictions that we bring to the data. Predictions 1–3 relate to which kinds of health signals trigger vaccination dropout, whereas 4 and 5 concern who responds more.

1. *Attribution.* Health events following vaccination reduce uptake of subsequent doses only if the symptoms can plausibly be attributed to the vaccine.
2. *Timing.* A shorter time from vaccination to symptom onset increases the effect on vaccination uptake.
3. *Salience and cues.* Effects are larger when medical professionals or public discourse reinforce the perceived link between the symptoms and the vaccine.
4. *Benefit gradient.* Individuals with lower predicted benefit from the vaccine respond more to health events.
5. *Hesitancy amplification.* Individuals who delay their first vaccination reduce their uptake of the next dose more in response to a health event.

The last two predictions share a common logic. Among those who took the first dose, individuals with lower $B(x_i)$ and/or higher t_i have lower u_i—closer to the threshold $u_i = 0$—and are therefore more likely to be tipped into dropping out upon suffering a negative health event.

4 Data

4.1 Administrative Records

We combine data from several Swedish administrative sources covering the entire Swedish population. We access the data through the Swedish Register-based Research Program on COVID-19 (SWECOV) and permission to use it is obtained from Sweden's Ethical Review Authority (permit numbers 2021-02225, 2022-013550-02, 2022-06118-02, and 2024-02342-02).

COVID-19 vaccination The Public Health Agency provides us with individual-level data on all COVID-19 vaccinations in Sweden between December 2020 and March 2023. These records include the administration

[5]Formally, for any u, u': $u' > u \implies G(\cdot|u)$ stochastically dominates $G(\cdot|u')$.

date for each dose as well as the vaccine brand and manufacturer, allowing us to capture vaccine hesitancy both through vaccination discontinuation and through vaccine brand switching.

Deaths, diagnoses & drugs We use the Swedish Death Register to identify individuals who died, either due to their blood clot conditions or from other causes. Furthermore, we exploit data on healthcare visits from the patient register administered by the National Board of Health and Welfare. This data includes detailed information on all specialist care visits within the period of 2005–2022, including the date and diagnoses. We use this data to identify cases of blood clot, as well as other conditions, based on the ICD-10-SE diagnosis classification.

Socioeconomic & demographic characteristics We access information from Statistics Sweden's registers on individual socioeconomic and demographic characteristics such as age, sex, region, and income. In addition, we rely on three sources to create proxies for health literacy. (1) From the military archive, we gather data from military tests that assess cognitive ability. (2) We use data on educational attainment and occupations to identify individuals who have a university degree or (3) who have a healthcare practitioner in their family.

Reported Side Effects We draw on individual-level records of suspected adverse drug reactions (ADRs) reported to Sweden's national spontaneous-reporting system throughout the COVID-19 pandemic, obtained from the Swedish Medical Products Agency between 2005 and 2022. Each entry contains the date of onset, the suspected drug link (ATC drug code and brand name), and the type of symptom, allowing us to identify individuals who had a blood clot reported as a suspected adverse event following COVID-19 vaccination. We restrict the attention to suspected side effects reported for blood clot-related symptoms as described in Appendix D.2.

4.2 Variable definitions and sample restrictions

Blood clots We are interested in the effect of experiencing a blood clot after vaccination on vaccine hesitancy. To this end, we first restrict our sample to individuals who have received at least one COVID-19 vaccine. This leaves us with individuals who have a basic willingness to get vaccinated. We further restrict the sample to adult individuals who were alive by 2023. We use a broad definition of blood clot diagnoses, including conditions that were highly unlikely to have been caused by the vaccine. Instead, the set of diagnoses included is meant to reflect blood clot related conditions that the individual might have inferred to be due to the vaccine. This allows us to draw conclusions about the generalization of the specific severe side effect of VITT to other blood clot conditions. We consider individuals with a blood clot diagnosis within 42 days of their first or second COVID-19 vaccination. The interval between the first and second dose was generally about six weeks. As a result, cases with longer intervals are likely a selected group of individuals who may already be hesitant, making the average effect in this group harder to interpret. Moreover, vaccine-induced blood clots were typically observed within six weeks, making the cases we study more comparable to the suspected adverse event. We also exclude individuals who developed a blood clot

after doses beyond the second. Decisions about fourth or later doses are less central to the vaccination schedule than decisions about the first three doses.

We exclude individuals diagnosed with a blood clot in the five years preceding the COVID-19 pandemic. This allows our treatment to more closely mirror the actual VITT cases, which were unlikely to occur disproportionately among individuals with a prior history of blood clots. A back of the envelop calculation suggests that nine individuals in Sweden developed VITT from AstraZeneca. We account for potentially true cases of VITT by excluding a handful of individuals from the dataset that had clinically diagnosed blood clots along with low blood platelet levels. This leaves us with a total of 8,497 individuals who developed a blood clot soon after COVID-19 vaccination.

Table 1: Distribution of diagnoses among individuals developing a blood clot following vaccination

Diagnosis (ICD-10)	**Share (%)**	**Number of individuals**	**Mortality rate (% deceased before 2023)**
Acute myocardial infarction (I21)	30.2	2,563	22.0
Cerebral infarction or stroke (I63)	27.8	2,358	28.2
Phlebitis and thrombophlebitis (I80)	21.4	1,822	14.9
Pulmonary embolism (I26)	15.7	1,338	33.5
Other venous embolism and thrombosis (I82)	2.4	204	29.8
Arterial embolism and thrombosis (I74)	1.9	160	31.0
Portal vein thrombosis (I81)	0.6	52	49.0
Total	100	8,497	100

Notes: This table displays the prevalence of blood clot indications in the sample of individuals living in Sweden in 2021–2023 (columns 1 and 2). Diagnoses are included if recorded within 42 days of the first or second COVID-19 vaccination. Column 3 presents a measure of all-cause mortality for each diagnosis code, referring to fatal cases within each category.

In Table 1 we display the diagnoses that we consider in our definition of blood clots.[6] The most common diagnoses in Table 1—myocardial infarction (I21), stroke (I63), phlebitis/thrombophlebitis (I80), and pulmonary embolism (I26)—are strongly related to underlying cardiovascular risk and comorbidities (including age as well as chronic and lifestyle-related conditions), and they do not share the hallmark presentation of vaccine-induced thrombosis at unusual sites together with thrombocytopenia. Nevertheless, patients may find it plausible that their diagnosis is vaccine-related because these events all involve the obstruction of blood flow. The displayed mortality rates refer to the share of individuals that are excluded from the sample because they die at some point after developing a blood clot and before 2023.[7] As evident from column (4) in Table 1, mortality differs markedly across diagnoses, but the post-diagnosis mortality observed within 1–2 years indicates that these events are severe in clinical terms.

In Figure 1 we display the number of first-time blood clot diagnoses per month between 2019 and 2023. There was no aggregate increase in blood clot incidence when COVID-19 vaccination was rolled out at

[6]Throughout the remainder of the paper, we refer to the diagnoses listed in Table 1 collectively as "blood clots" for expositional convenience, although some of these diagnoses are not blood clots in a strict clinical sense.

[7]We abstain from using available information on the cause of death to restrict our sample, because the measure disregards that the primary cause of death may instead be a comorbidity of the blood clot condition.

scale. Blood clots are common: each year, about 0.5% of individuals receive a first-time diagnosis. This high baseline incidence implies that, even at the time, clinicians would expect many post-vaccination blood clots to occur for reasons unrelated to the vaccine. Out of the 22,957 individuals who developed a blood clot after COVID-19 vaccination, 443 individuals in addition had serious adverse events for blood clot-like symptoms reported by medical professionals. The average age among individuals developing a blood clot after COVID-19 vaccination is 74 years and split equally by gender.

Ex ante, individuals could nevertheless worry that rare vaccine-related cases were underdetected or not yet reported. By March 2021, however, millions of AstraZeneca doses had already been administered worldwide and only a small number of suspected cases had been identified. Contemporaneous evidence thus suggested that the probability of a newly diagnosed blood clot being vaccine-induced was very low. We take the view that individuals did not fully incorporate this evidence, and it is precisely this discrepancy between best available knowledge and individual beliefs that creates scope for personal experiences to shape vaccination behavior.

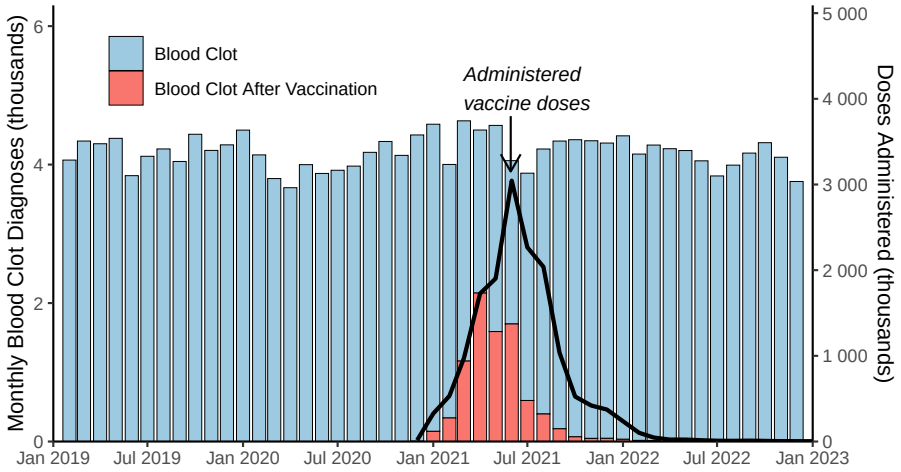

Figure 1: Timeline of Blood Clot Diagnoses

Notes: This figure displays first time blood clot diagnoses across time between 2019 and 2022 (left axis). The blue bars represent the total count of diagnoses. The red bars represents individuals that developed blood clot within six weeks after receiving either their first or second COVID-19 vaccine dose and hence constitute our main treatment group. The black line represents the monthly number of first and second administered COVID-19 vaccine doses (right axis).

Other conditions Apart from blood clots, we construct a group of *physical injuries*, consisting of all ICD-10 S-codes (Chapter XIX: Injury, poisoning and certain other consequences of external causes), which capture acute physical injuries to head, trunk and limbs. These diagnoses are unlikely to be attributed to the COVID-19 vaccine. We expect any effects among individuals with these conditions to, if anything, reflect an incapacitation effect where for example the physical restrictions from breaking a leg makes it more difficult to get another dose of COVID-19. We also construct a set of *other acute conditions* consisting of other acute conditions such as infectious diseases, kidney failure and appendicitis. See Appendix D.1 for a more detailed description of how we select diagnoses for this category. These conditions are broadly similar in severity to blood clots and it is plausible that a layperson would have attributed the condition to the vaccine to a similar

extent as blood clots, had blood clots not been made especially salient in the media.

In summary, we distinguish between three sets of diagnoses: (i) blood clots, which may be conflated with VITT; (ii) physical conditions, which are primarily incapacitating the range of movement and unlikely to be attributed to the vaccine; and (iii) other severe acute conditions, which could plausibly be attributed to the vaccine but where there was no media coverage.

Outcome variables We exploit the information on vaccination doses to define an outcome variable for subsequent vaccination, equal to one if an individual takes at least one more COVID-19 dose before March 2023, which marks the end of our data but also the end of the acute phase of the COVID-19 pandemic, independent of the brand of that next dose. We also study *switching*, equal to one if an individual continues vaccination but with a different brand than the brand that preceded the blood clot.

5 Method

Challenges to Identification & Matching Our empirical strategy compares immunization outcomes among individuals who developed a blood clot shortly after receiving a COVID-19 vaccine to those who received the same vaccine but did not develop a blood clot. The identifying assumption is that conditional on observable characteristics, blood clot diagnosis is orthogonal to the latent propensity to continue vaccinating. Two threats to our identification strategy stand out. First, individuals may differ in their probability of developing a blood clot in ways that correlate with their inclination to get vaccinated. For example, individuals with comorbidities may be more likely to develop blood clots and also benefit more from continued vaccination. Second, individuals who suffer a blood clot may differ in their propensity to get a formal diagnosis. For example, those who live far from a healthcare facility or distrust the healthcare system may be less likely to receive a diagnosis and less likely to continue vaccinating. In both examples, the treatment group would be positively selected on vaccination propensity, biasing our estimates toward zero. The net direction of bias is, however, ambiguous, since other selection pressures could work in the other direction.

To address these endogeneity concerns, we deploy a two-step matching procedure. The key idea is to match individuals with identical COVID-19 vaccination histories and similar ex-ante probabilities of a blood clot diagnosis. Concretely, each treated individual is first matched to a pool of potential controls consisting of everyone who had taken the exact same vaccine brands in the exact same months and years, and had not yet taken their next dose by the time of the treated individual's diagnosis.[8] By matching on vaccine brand together with vaccination timing, we address what is perhaps the most important selection concern: individuals who are more worried about side effects selecting into vaccines perceived as safer and receive their doses later. Within the matched control pool, we select the one nearest untreated neighbor based on propensity scores estimated on the full sample using XGBoost. Beyond remaining agnostic about functional form, this tree-based method captures nonlinearities and interactions that would require many additional terms in a parametric logit model. The scores capture blood-clot risk using medical histories, including the

[8] As an example, consider a treated individual who is diagnosed with a blood clot 28 days after her second vaccination dose. Her control pool consists of all individuals who took the same brands as her for their first and second doses, in the same calendar months. Furthermore, the pool is limited to individuals who had not yet received their third vaccine dose within 28 days after their second dose.

number of drugs taken and healthcare visits, together with socioeconomic and demographic characteristics such as years of schooling and income. This step ensures that each pair of treated and control individuals faced a similar ex-ante probability of being diagnosed with a blood clot. Section C explains the procedure in greater detail.

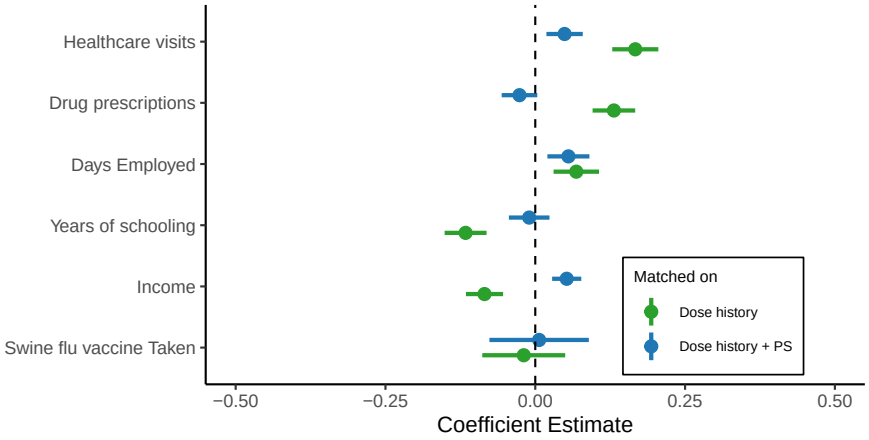

Figure 2: Balance in Pre–Treatment Characteristics

Notes: This figure displays results from univariate regressions where the standardized covariates is regressed on treatment (diagnosed with blood clot). In the "Dose history" sample we match on (i) vaccine brand(s) received, (ii) calender month of vaccination. Each treated individual is, at random, matched to one control individual who had not yet taken their next dose by the time the treated individual is diagnosed with the blood clot. "Dose history + PS-matching" is instead based on a sample where to each treated individual we match a control individual based on propensity scores as described in Appendix C. Horizontal lines show 95% confidence intervals based on standard errors clustered at the match-pair level.

In Figure 2, we display differences in a selection of predetermined characteristics between treatment and control groups, before and after matching on propensity scores. Matching solely on vaccine dose history leaves residual imbalance in terms of individual health characteristics as measured through healthcare visits and drug prescriptions. Notably, while there is little difference in socioeconomic characteristics, individuals in the treatment group appear to be sicker as measured by the number of drug prescriptions and healthcare visits. To approximate conditional independence, we match individuals on characteristics that may jointly predict both vaccine hesitancy and blood clot onset. The propensity score matching does a good job at achieving balance, although some residual imbalance remains. A concern is that latent vaccine hesitancy may only be weakly correlated with standard observable health and socioeconomic characteristics. Fortunately, we access data on swine flu vaccines taken during the 2009–2010 swine flu pandemic for a subset of our sample. We observe no statistically significant differences in swine flu vaccination status neither before or after matching. In other words, developing a blood clot after receiving a COVID-19 vaccine is not associated with a reluctance to take the swine flu vaccine 11 years prior to the COVID-19 pandemic.

6 Results

6.1 Which conditions cause vaccine dropout?

In Figure 3 we display our main results for the different blood clots along with estimates for physical conditions and other severe acute conditions.

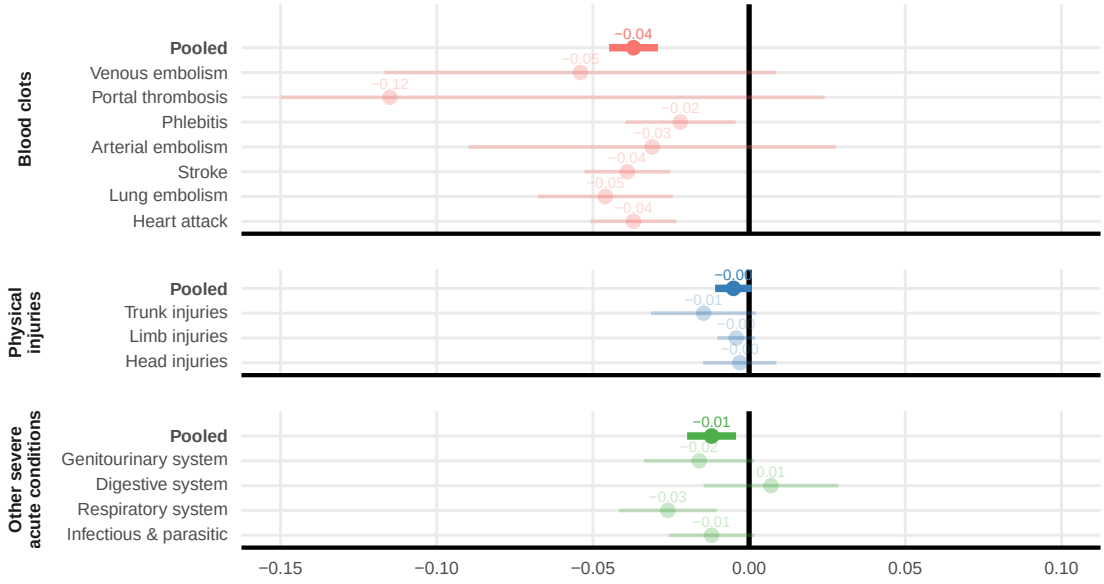

Figure 3: Vaccine Hesitancy and Media Exposure

Notes: This figure displays coefficient estimates of regressions run on a matched sample where the dependent variable is equal to 1 if an individual takes an additional COVID-19 vaccine doses and 0 otherwise. We consider three types of conditions: (i) Blood clot disorders that were unlikely to have been vaccine-induced but plausibly believed to have been induced by the COVID-19 vaccine; (ii) Physical conditions that were highly unlikely to have been adverse events and where the individual is unlikely to infer that these were adverse events; (iii) other severe acute conditions that could ex ante plausibly be attributed to vaccination but were not salient in public discussion and were unlikely to be vaccine-induced. Vertical lines show 95% confidence intervals based on standard errors clustered at the match-pair level.

As shown in the top panel, we find an average reduction in subsequent vaccine uptake of 4 percentage points among individuals who experienced a blood clot soon after vaccination.[9] There is some heterogeneity across specific blood clot diagnoses, although our data does not allow for sharp conclusions on this point. Most remarkably, the diagnoses most similar to VITT in terms of clot location (venous embolism and portal thrombosis) return the largest estimates.

We also find negative effects for other severe acute conditions (bottom panel), for instance relating to the genitourinary system or the respiratory system, albeit smaller in magnitude and mostly statistically

[9] In an alternative specification, we construct a control group where control units are drawn from a donor pool of individuals who developed blood clot *before* the COVID-19 pandemic. This approach directly creates comparable units in terms of propensity to develop blood clot. However, this group of individuals is smaller, rendering it less feasible to target imbalances in other characteristics such as socioeconomic characteristics. Furthermore, they differ from the main candidate pool in that they selected into taking the first COVID-19 dose despite developing blood clot in the past. Either way, the approach yields identical results for blood clot diagnoses, see Figure B4.

insignificant. This contrast in results is unlikely to reflect differences to blood clots in terms of severity: as shown in Table D1, these conditions are also severe and associated with high mortality rates. Instead, the small effects are consistent with misattribution: some individuals may attribute acute conditions occurring soon after vaccination to the vaccine, even absent media coverage on their particular condition.

We do not detect an effect for individuals experiencing physical injuries, such as a foot injury. This is in line with the notion that individuals cannot plausibly connect their diagnosis to the vaccine and therefore do not adjust their beliefs towards the vaccine.

Taken together, we find evidence of a small degree of misattribution stemming from developing an unrelated disease soon after vaccination in itself but the media coverage of rare and specific blood clots amplified hesitancy among individuals who experienced similar events shortly after vaccination.

Spillover to family members To capture the overall vaccine hesitancy response to an additional blood clot case, regardless of whether it arises for the affected individual or for others, we study spillovers within personal networks. In Table B1, we report the estimates for partners, siblings, and children of the diseased individual. The outcome variable is defined analogously to our previous approach, equaling 1 if a network member takes another vaccine dose after the focal individual within their family is diagnosed with a blood clot, and 0 otherwise. We report our estimates separately for the case where the focal individual died or lived on after experiencing a blood clot. Overall, the estimated effects are negative with the magnitude being about 30–60% of the magnitude for diseased individuals themselves, although the estimates for children are virtually null. Effects among siblings, who share genes with the diseased individuals, are not larger than those for partners. This pattern suggests that the responses we document are not solely driven by individuals updating beliefs about their own predisposition to side effects, instead individuals use experiences to infer the common risk of severe side effects from the next vaccine dose. These results relate to Miltner and Riberth (2026) which finds that close family members of individuals developing narcolepsy after receiving the swine-flu vaccine were significantly less likely to be vaccinated during the COVID-19 vaccination campaign.

Accounting for the spillovers to family members, the point estimates in Table B1 suggest that each blood clot case is associated with 0.08 individuals dropping out of the vaccination schedule.[10] In other words, about half of the effect on stopping vaccination comes from the affected individual, and the other half comes from family members not continuing vaccination.

6.2 Which signals matter most?

Heterogeneity by time until blood clot Our explanation for why individuals become more hesitant after experiencing a blood clot is that they form an assessment of the risk of developing side effects from the next dose. This assessment in turn depends on their perceived likelihood that the blood clot they developed was induced by the previous vaccine dose. On this view, the perceived link should be stronger for individuals who developed a blood clot closer in time to vaccination. If, instead, dropout were driven by incapacitation

[10]The number comes from multiplying the estimated effect sizes in Table B1 with the size of each network among the treated units (2.33 children, 2.1 siblings and 1 partner). For comparability to the main estimate, we focus on the case where the focal individual survives.

due to sickness, we would expect the opposite pattern: blood clots would have a greater effect the closer they occur to the date of the planned *next* dose. We test these hypotheses in Figure 4, where we split the sample by the number of weeks elapsed between the COVID-19 vaccination and the diagnosis of the blood clot. For comparison, we also display the corresponding estimates for other acute conditions.

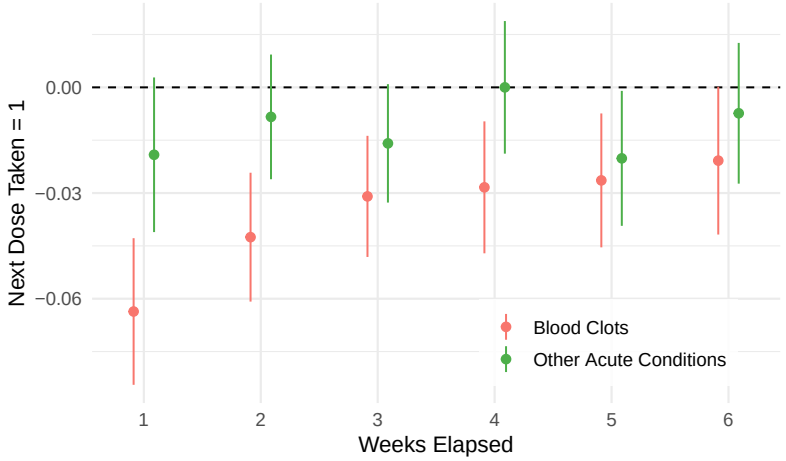

Figure 4: Heterogeneity by Time Elapsed Between Vaccine and Blood Clot

Notes: This figure displays coefficient estimates of regressions run on a matched sample where the dependent variable is equal to 1 if an individual takes an additional COVID-19 vaccine doses and 0 otherwise. The sample is split by the number of weeks elapsed between the last COVID-19 vaccine dose and timing of diagnosis. Error bars represent 95% confidence intervals based on standard errors clustered at the match-pair level.

The results indicate that the dropout effect from blood clots decreases sharply with the length of the latency period. This pattern is in line with the hypothesis that dropout is due to causal attribution, and hard to reconcile with alternative explanations. The gradient is absent for other acute conditions, indicating that the small aggregate effect from these conditions may reflect general disruption from illness more than attribution to the vaccine. In Figure B5 we display the corresponding results for the decision to switch vaccine brand, showing qualitatively similar though less stark results. Taken together, individuals appear to base their perceived risk of side effects and hence the decision to continue vaccination on the perceived probability that the blood clot was indeed vaccine induced.

The role of doctors' communication Our results suggest that reports of suspected side effects led to media coverage which caused individuals to attribute unrelated blood clots to the vaccine. To shed further light on the role of external cues, we consider effects among a subset of 394 individuals who, on top of developing a blood clot after COVID-19 vaccination, had the event reported as a suspected side effect by a healthcare professional. The premise is that doctors communicate their decision to report to the patient, and that this communication provides a cue akin to media coverage—signaling to the individual that the blood clot is vaccine-induced. We find significantly larger effects in this subset with an average effect of 16 percentage points compared to the effect of solely developing a blood clot (4 percentage points). Our interpretation is that doctors report it as a side effect and in doing so cues the patient perception that the condition was vaccine-induced, increasing the perceived probability that the individual will develop side

effects from subsequent doses as well.[11]

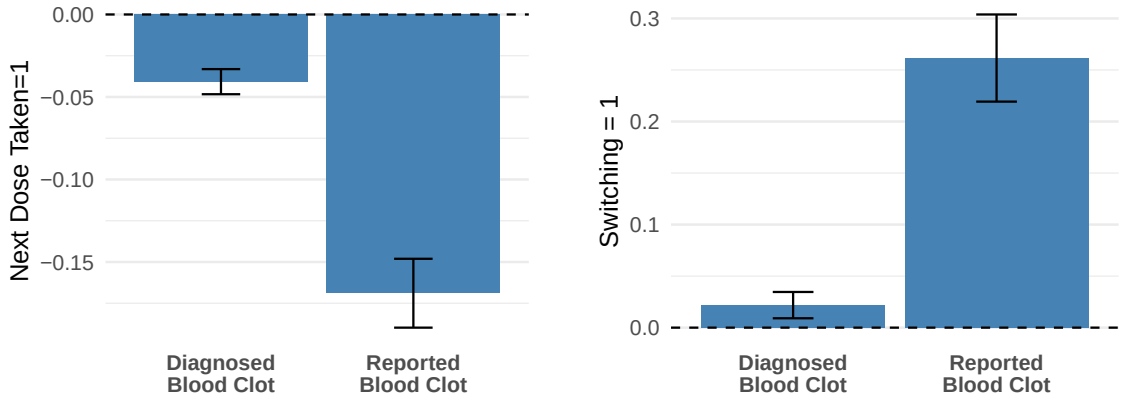

(a) Continuing Vaccination

(b) Switching Vaccines

Figure 5: Role of doctors' communication

Notes: This figure reports coefficient estimates from regressions estimated on the matched sample. Diagnosed BC indicates individuals diagnosed with a blood clot within six weeks of receiving a COVID-19 vaccine (or their matched controls), while Reported BC further restricts the sample to cases reported as suspected adverse events by a healthcare professional. **Panel (a)** uses a binary indicator for continued COVID-19 vaccination as the dependent variable. **Panel (b)** uses a binary indicator for switching vaccine brand as the dependent variable. Error bars represent 95% confidence intervals based on standard errors clustered at the match-pair level.

These results speak to an important policy trade-off. While doctors need to report suspected side effects to detect new ones, the act of reporting side effects also causes individuals to become more hesitant, despite the fact that the experienced health event was unlikely to have been a side effect.

Heterogeneity across vaccine brands We study heterogeneity across vaccine brands to assess whether individuals' responses were narrowly concentrated on AstraZeneca—the vaccine linked to VITT—or whether concerns spilled over to other brands as well. Spillovers would be consistent with imperfect attribution and "generalization" in risk beliefs, and would imply a broader misattribution that extends not only to non-VITT blood clot conditions, but also to vaccine brands that were not implicated. We present the results of this exercise in Figure 6.

For the decision to take the next dose, there is little heterogeneity across brands, suggesting that individuals who drop out of the vaccination schedule do not reflect on which vaccine might have caused the adverse event.

Many individuals, however, respond by continuing vaccination but switching brand. Experiencing a blood clot after vaccination with AstraZeneca or Moderna significantly increases the probability of switching, with the largest effects following vaccination with AstraZeneca. The control-group switching rates are 12% for Pfizer, 26% for Moderna, and 48% for AstraZeneca. We do not detect comparable switching responses from blood clots following Pfizer. This is consistent with Pfizer being perceived as the safest vaccine: while

[11] One might expect stronger responses among women because early reporting around the AstraZeneca–blood clot episode often emphasized cases in younger women, potentially making the side effect narrative more salient for females. Figure B6 shows no evidence of such gender heterogeneity in either continuation or switching behavior.

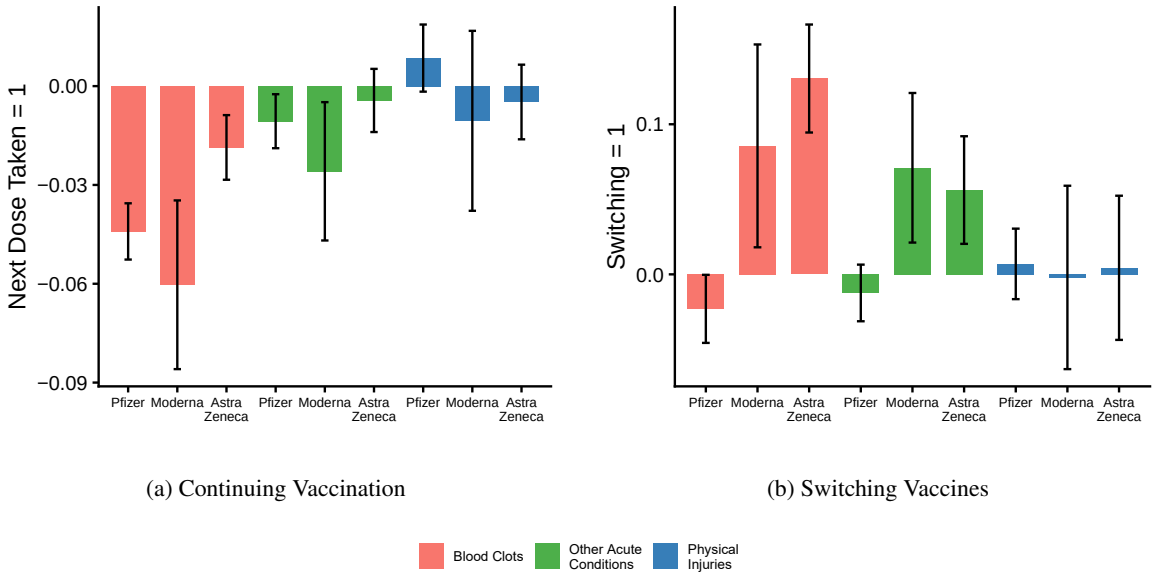

(a) Continuing Vaccination

(b) Switching Vaccines

Figure 6: Brand Heterogeneity

Notes: Sample split up by the brand the individual received prior to developing a blood clot. **Panel (a)** uses an indicator for continued COVID-19 vaccination as the dependent variable. **Panel (b)** uses an indicator for switching vaccine brand as the dependent variable. Error bars represent 95% confidence intervals based on standard errors clustered at the match-pair level.

abstention effects look similar, the incentive to switch away from Pfizer is comparatively smaller.

Finally, switching effects are larger for AstraZeneca in the blood clot analysis than for the other acute conditions, whereas the difference is smaller for Moderna. This pattern suggests that media exposure played a role in amplifying hesitancy and switching around AstraZeneca.

To account for compositional differences in who received the three brands, we reweight observations in the Pfizer and Moderna groups by age, risk group status and whether or not they have a doctor in their family to match the AstraZeneca group following DiNardo et al. (1996). We display these results in Figure B3. While the results are overall similar, it is worth noting that the effects are generally smaller—especially for the decision to take the next dose among individuals who experience a blood clot, and for the decision to switch vaccines after developing a blood clot following Moderna. This likely reflects that the AstraZeneca group is older and at higher risk of severe COVID-19.

6.3 Who responds most?

Heterogeneity in benefits of vaccination Abstaining from vaccination entails an excess risk of falling ill from COVID-19 and experiencing a severe course of disease. Do individuals at higher risk of severe COVID-19 respond less to exposure to side effects? In Figure 7, we use two proxies for the perceived benefits of vaccination: age and the time elapsed between when the vaccine became available to an individual and the timing of the first dose. Older individuals face a higher risk of severe COVID-19 and therefore benefit more from vaccination than younger individuals. A similar pattern holds for time since availability. Individuals who take the vaccine shortly after it becomes available are likely those who perceive vaccination as more beneficial and are therefore less likely to stop vaccination after developing a blood clot. Following the same

logic, individuals who delay vaccination have a baseline hesitancy, and are the ones for whom a shock to perceived costs is more likely to tilt the decision toward abstaining from future vaccination.

We find that older individuals, and those who took the dose preceding the blood clot diagnosis earlier in time, react less than younger individuals or those who took their first dose later.

Although media may have emphasized a higher likelihood of developing VITT for younger individuals, combining panels (a) and (b) of Figure 7 reveals a benefit–cost trade-off: perceived side-effect risk is similar across individuals experiencing a blood clot, but perceived benefits vary, generating heterogeneity in the decision to stop vaccinating. Finally, there is a substantially smaller gradient for other acute conditions as compared to blood clot diagnoses. This suggests that these individuals do not trade off costs and benefits in the same way as those experiencing blood clots.

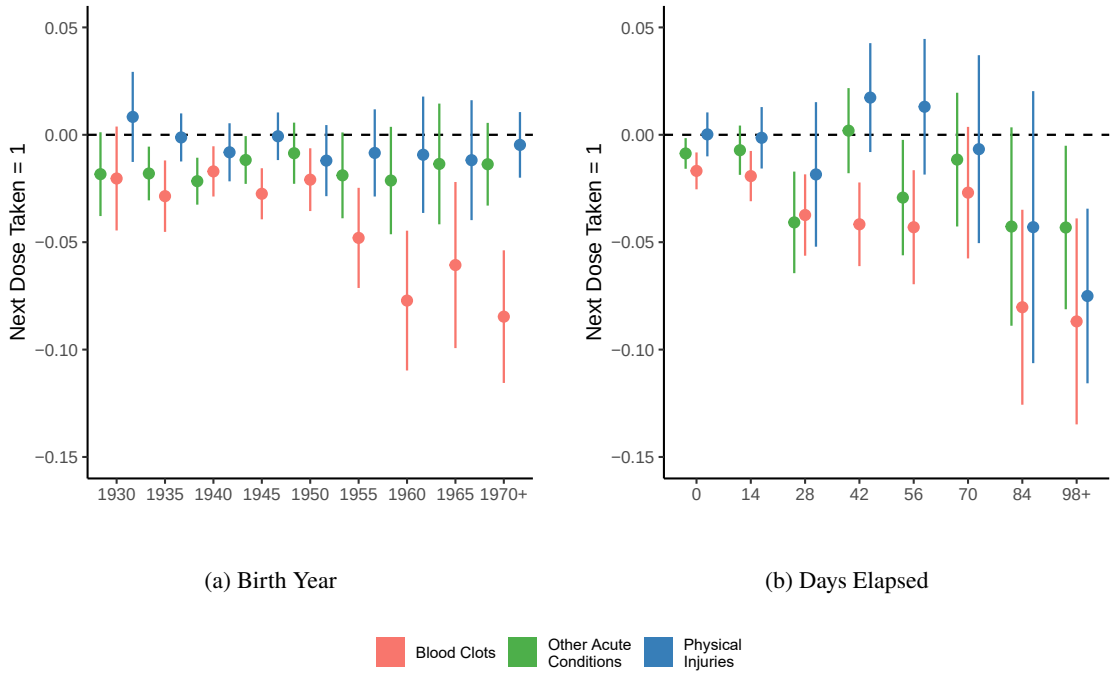

(a) Birth Year

(b) Days Elapsed

Figure 7: Eliciting the perceived benefits of vaccination

Notes: This figure displays coefficient estimates of regressions run on a matched sample where the dependent variable is equal to 1 if an individual takes an additional COVID-19 vaccine doses and 0 otherwise. **Panel (a)**: Sample split up by birth year. **Panel (b)** Sample split up by days elapsed since first date the COVID-19 vaccine is available. Date of first availability is defined in Appendix D.3. Error bars represent 95% confidence intervals based on standard errors clustered at the match-pair level.

Heterogeneity in health literacy Finally, we study heterogeneity by personal health literacy. We are interested in how health literacy interacts with how individuals understand and interpret cues from media coverage of VITT. For instance, having a doctor in the family may make it easier to understand the lack of a clinical association between common blood clot conditions and the rare cases of VITT. Figure 8 presents our main estimates by high and low health literacy using three alternative measures of health literacy.

We find no statistically significant heterogeneity with respect to health literacy. If anything, the coefficients suggest that the blood-clot signal leads individuals with higher health literacy to react less to experiencing a blood clot. Taken together, while health literacy matters for vaccination uptake, our findings suggest

that the interpretation of personal experiences is remarkably similar across individuals. Even though one might expect individuals with higher health literacy to discount spurious signals, both high- and low-literacy individuals appear to attribute such health events to the vaccine in similar ways.

This finding is consistent with Miltner and Riberth (2026), who show that—irrespective of health literacy—exposure to vaccine-induced narcolepsy following the swine flu vaccine led to substantial vaccine drop-out during the COVID-19 pandemic.

In the broader literature, our results resonate with Malmendier et al. (2021), who show that past experiences of high inflation shape inflation expectations and voting behavior on monetary policy even among highly educated and well-informed members of the U.S. Federal Open Market Committee.

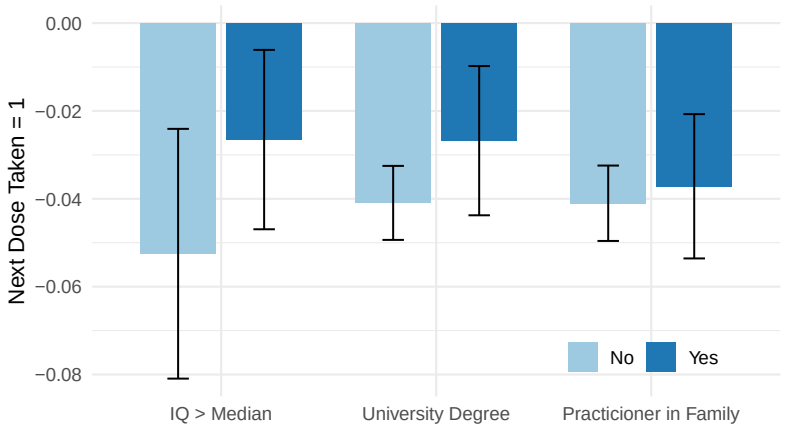

Figure 8: Heterogeneity in Health Literacy

Notes: This figure displays coefficient estimates of regressions run on a matched sample where the dependent variable is equal to 1 if an individual takes an additional COVID-19 vaccine doses and 0 otherwise. We consider three proxies for health literacy. *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. Error bars represent 95% confidence intervals based on standard errors clustered at the match-pair level.

7 Conclusions

This paper shows that experiencing an adverse health event after vaccination can lead individuals to falsely attribute symptoms to the vaccine. Informational cues, such as salient media attention surrounding side effects or a clinician's decision to file a side-effect report, prompt this response. In our setting, experiencing a blood clot—clinically distinct from the rare vaccine-induced condition—shortly after COVID-19 vaccination reduces the probability of continuing the vaccination schedule. This pattern is consistent with individuals generalizing from the most salient side-effect narrative to their own symptoms. By contrast, we detect no meaningful change in vaccination behavior after conditions that present differently from the blood clots, highlighting the role of media coverage in the interpretation of symptoms following vaccination.

Theoretically, our results highlight how people generalize information across health domains: effects

spill over not only to blood clot diagnoses that differ from the specific adverse event of concern, but also across vaccine brands, even when only one brand is linked to elevated risk. Taken together, this points to substantial misattribution, with individuals extrapolating from a rare, brand-specific adverse event to other blood clot diagnoses and to vaccines for which no elevated risk was suspected.

From a policy perspective, the results suggest that the communication about side effects carries risk for vaccine campaigns. The costs arise as soon as news about a potential side effect breaks: the withdrawal of the AstraZeneca vaccine shortly thereafter did not prevent a persistent increase in hesitancy among those who experienced blood clots in the same period. The evidence on clinician side-effect reporting also points to a similar trade-off. Reporting is essential for detecting genuine adverse reactions, yet the act of reporting may unintentionally reinforce patients' beliefs that their condition was vaccine-induced. Importantly, COVID-19 vaccination in Sweden was not broadly mandated. In settings where policies create stronger incentives to vaccinate, for example through workplace mandates or vaccine passports, the behavioral responses documented here may lead to larger effects on vaccine hesitancy, as such policies may induce vaccination among individuals with higher underlying hesitancy who are more likely to react strongly to adverse health events.

The practical scope of our findings may be substantial. In the MMR–autism case, a highly salient medical claim was later retracted, yet many people did not fully revise their beliefs. Instead, vaccine confidence remained lower for years following the retraction (Miltner and Riberth, 2026). Evidence of such durable effects, even after the underlying claim was discredited, suggests that the AstraZeneca–blood clot episode may likewise generate persistent effects.

Public debates about vaccine side effects are not uncommon, but which risks become "the" public story can be somewhat arbitrary. For example, myocarditis and pericarditis were identified as very rare adverse events following the mRNA vaccines, concentrated among young males, yet this risk did not generate the same sustained debate as the AstraZeneca–blood clot episode. This suggests that media, clinicians, and other stakeholders shape which risks are amplified in public perception, and that salient outcomes such as blood clots may be especially likely to dominate attention. The lesson for risk communication is that when attention is driven by salience as much as by incidence, some risks can be amplified out of proportion while other clinically important risks remain underweighted. Future research should therefore focus on how these narratives form and spread in increasingly digital information environments, and on how institutions can communicate safety signals in ways that preserve trust while keeping perceived risks aligned with the underlying evidence.

References

Agosti, F., V. Toffolutti, N. Cavalli, S. Nivakoski, M. Mascherini, and A. Aassve (2022). Information and vaccine hesitancy: Evidence from the early stage of the vaccine roll-out in 28 european countries. *PLOS One 17*(9), e0273555.

Asan, L., J. Kleine-Borgmann, B. Bozkurt, B. Frank, M. Köhrmann, C. Kleinschnitz, and U. Bingel (2024). Media coverage of covid-19 vaccination-associated cerebral venous sinus thrombosis was followed by a surge in emergency presentations due to headache–observations from a university hospital in germany. *Frontiers in Psychiatry 15*, 1378472.

Ashraf, N., G. Bryan, A. Delfino, E. Holmes, L. Iacovone, C. Meyer, and A. Pope (2024). Learning to see a world of opportunities. *Working Paper. Accessed November 22, 2024.*.

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

Deiana, C., A. Geraci, G. Mazzarella, and F. Sabatini (2022). Perceived risk and vaccine hesitancy: Quasi-experimental evidence from Italy. *Health Economics 31*(6), 1266–1275.

DiNardo, J., N. M. Fortin, and T. Lemieux (1996). Labor market institutions and the distribution of wages, 1973-1992: A semiparametric approach. *Econometrica 64*(5), 1001–1044.

Espín-Sánchez, J.-A., S. Gil-Guirado, and N. Ryan (2023, June). Praying for rain. NBER Working Paper 31411, National Bureau of Economic Research.

Franchini, M., G. M. Liumbruno, and M. Pezzo (2021). Covid-19 vaccine-associated immune thrombosis and thrombocytopenia (vitt): diagnostic and therapeutic recommendations for a new syndrome. *European Journal of Haematology 107*(2), 173–180.

Haas, J. W., F. L. Bender, S. Ballou, J. M. Kelley, M. Wilhelm, F. G. Miller, W. Rief, and T. J. Kaptchuk (2022). Frequency of adverse events in the placebo arms of covid-19 vaccine trials: a systematic review and meta-analysis. *JAMA network open 5*(1), e2143955–e2143955.

Johansson, M., L. Johansson, and M. Lind (2014). Incidence of venous thromboembolism in northern sweden (veins): a population-based study. *Thrombosis Journal 12*(1), 6.

Jones, S. P. (2020, April). Imperial college london yougov covid data hub. https://www.coviddatahub.com Online resource; last accessed 03/05/2026.

Katsoularis, I., O. Fonseca-Rodríguez, P. Farrington, H. Jerndal, E. H. Lundevaller, M. Sund, K. Lindmark, and A.-M. Fors Connolly (2022). Risks of deep vein thrombosis, pulmonary embolism, and bleeding after covid-19: nationwide self-controlled cases series and matched cohort study. *BMJ 377*.

Klok, F. A., M. Pai, M. V. Huisman, and M. Makris (2022). Vaccine-induced immune thrombotic thrombocytopenia. *The Lancet Haematology 9*(1), e73–e80.

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

Malmendier, U. and G. Tate (2005, November). Ceo overconfidence and corporate investment. *The Journal of Finance 60*(6), 2661–2700.

MHRA (2023). Coronavirus vaccine—summary of yellow card reporting (archived version updated 8 march 2023). https://www.gov.uk/government/publications/coronavirus-covid-19-vaccine-adverse-reactions/coronavirus-vaccine-summary-of-yellow-card-reporting--2 Online resource; last accessed 03/05/2026.

Miltner, S. and J. Riberth (2026). Adverse health events and vaccine hesitancy. *Unpublished Manuscript*. https://jonatanriberth.github.io/jmp/jmp.pdf.

Wang, J. J., L. Schönborn, T. E. Warkentin, L. Müller, T. Thiele, L. Ulm, U. Völker, S. Ameling, S. Franzenburg, L. Kaderali, et al. (2026). Adenoviral inciting antigen and somatic hypermutation in vitt. *New England Journal of Medicine 394*(7), 669–683.

Zhao, L., H. Tran, M. Risk, and G. Nair (2024). Risk of blood clots after covid-19 vaccination and infection: a risk-benefit analysis. *Research square*, rs–3.

Appendix A. Medical background of vaccine-induced blood clots

Vaccine-induced immune thrombotic thrombocytopenia (VITT), a form of thrombosis with thrombocytopenia syndrome (TTS), is an extremely rare adverse event that has typically been observed within 42 days of the viral vector COVID-19 vaccines of AstraZeneca and Johnson & Johnson. The mRNA vaccines Pfizer and Moderna have not been scientifically linked to an excess risk of VITT (Klok et al., 2022).

The clinical presentation of VITT depends on the precise location of the clot. The blood clots for VITT commonly form in unusual sites of the body (mainly cerebral or sinus vein thromboses), although different presentations are possible, i.e. a deep vein thrombosis.

Scientists only recently identified the cause for these strong but extremely rare physical reactions (Wang et al., 2026). In case of VITT, the body produces an immune responses against parts of the adenovirus, causing the hyperactivation and clumping of PF4. For this reason, a central diagnosis criteria for VITT is *low* blood platelet count, which is in stark contrast to other blood clot conditions. Medical practitioners are able to measure the count of blood platelets in laboratory testing. In the absence of a competing potential cause for the observation of low blood platelets, i.e. heparin administration, the diagnosis the diagnosis of VITT is confirmed by determining the date of the last COVID-19 vaccine.

The course of therapy is similar for patients with common blood clots and patients of vaccine-induced blood clots. The main target of therapy is the prevention of further growth in clots and formation of new clotting. This is commonly achieved by administrating anticoagulant medication, blood thinners, such as heparin injections or equivalent products for oral administration. Because of the distinct pathophysiology of VITT, heparin—often administrated in regular blood clots—could exacerbate the clotting. It is therefore sometimes exchanged for non-heparin based anticoagulants, which warrants medical personnel to test for low blood platelet count.

Appendix B. Additional Descriptive Statistics & Results

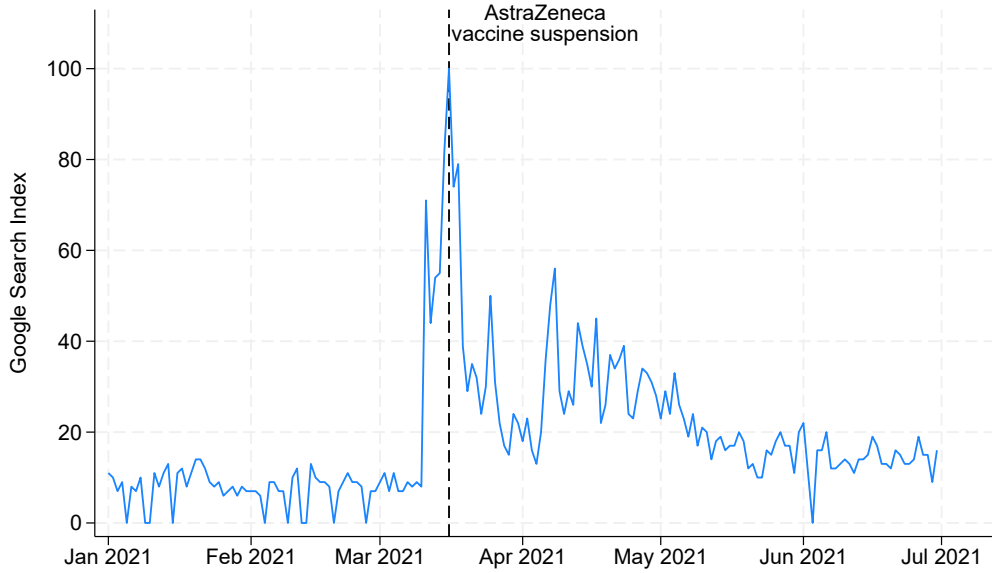

Figure B1: Timeline of Google Searches

Notes: This figure displays the daily count of google search index for Sweden using the keyword "blodpropp" between January 1st and August 31st 2021. The frequency of searches are indexed to the time of highest search volume (= 100). The vertical dashed line represents the timing of the AstraZeneca vaccine suspension in Sweden on March 16 2021. Data last accessed January 20, 2026.

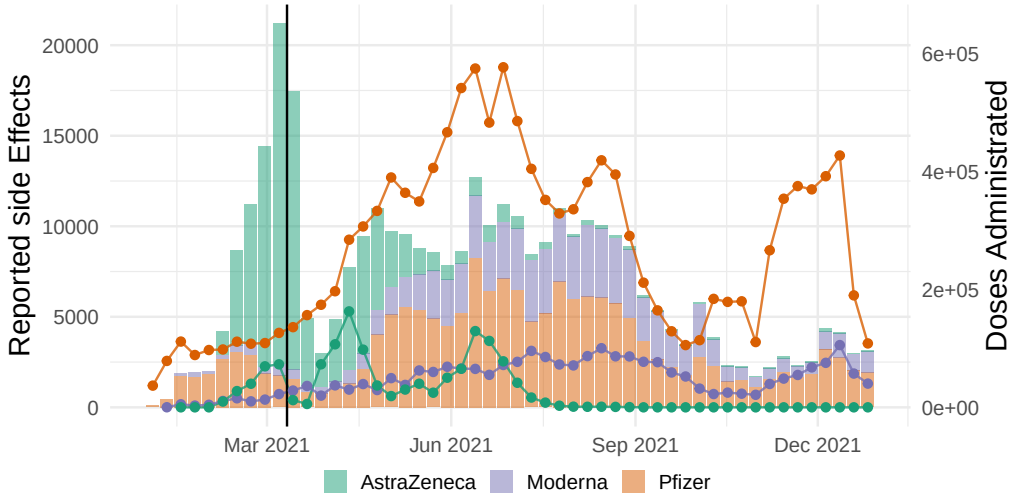

Figure B2: The Blood Clot Episode: Blood Clot Reports Over Time

Notes: This figure displays reported reported side effects for blood clot symptoms across time (bars). In addition, number of COVID-19 doses administered are displayed (lines). The black line represents the first report of a potential link between AstraZeneca and an excess risk of developing a blood clot.

Table B1: Effect of Blood Clot on Network Members' Vaccination Decisions

	Children	Partners	Siblings
Focal individual died = 0			
Treated	0.003	-0.019	-0.013
	(0.003)	(0.003)	(0.004)
n obs	84106	24590	59477
Focal individual died = 1			
Treated	0.005	-0.009	-0.018
	(0.004)	(0.006)	(0.008)
n obs	29190	6431	16155

Notes: This table reports estimates of the effect of exposure to a family member developing a blood clot shortly after vaccination on the vaccination uptake of network members. The sample of network members is restricted to adults. Partners are defined as either married or, if not married, cohabiting. Each column corresponds to a separate regression. The dependent variable is equal to 1 if an individual takes at least one COVID-19 vaccine dose after the family member is diagnosed with blood clot. Standard errors are clustered at the match-pair level and by the focal individual.

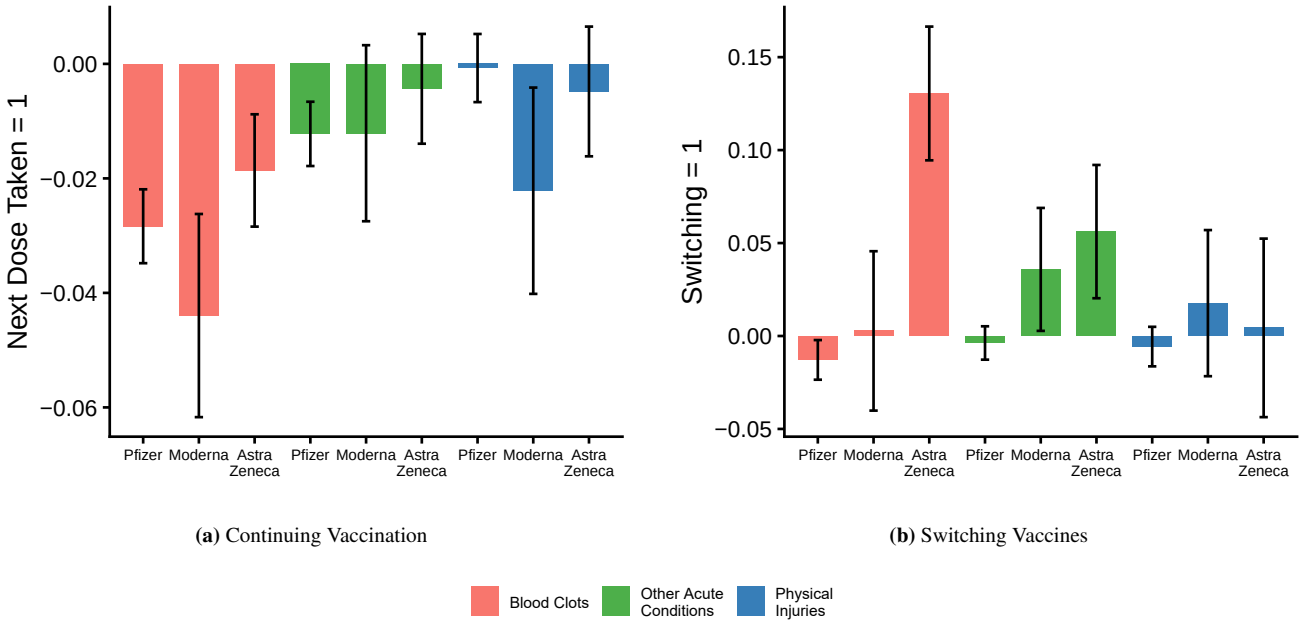

(a) Continuing Vaccination

(b) Switching Vaccines

Figure B3: Brand Heterogeneity, DFL reweighted

Notes: This figure displays coefficient estimates of regressions where the sample is split up by the brand that the individual received prior to developing a blood clot. **Panel (a)** uses an indicator for continued COVID-19 vaccination as the dependent variable. **Panel (b)** uses an indicator for switching vaccine brand as the dependent variable. Error bars represent 95% confidence intervals based on standard errors clustered at the match-pair level.

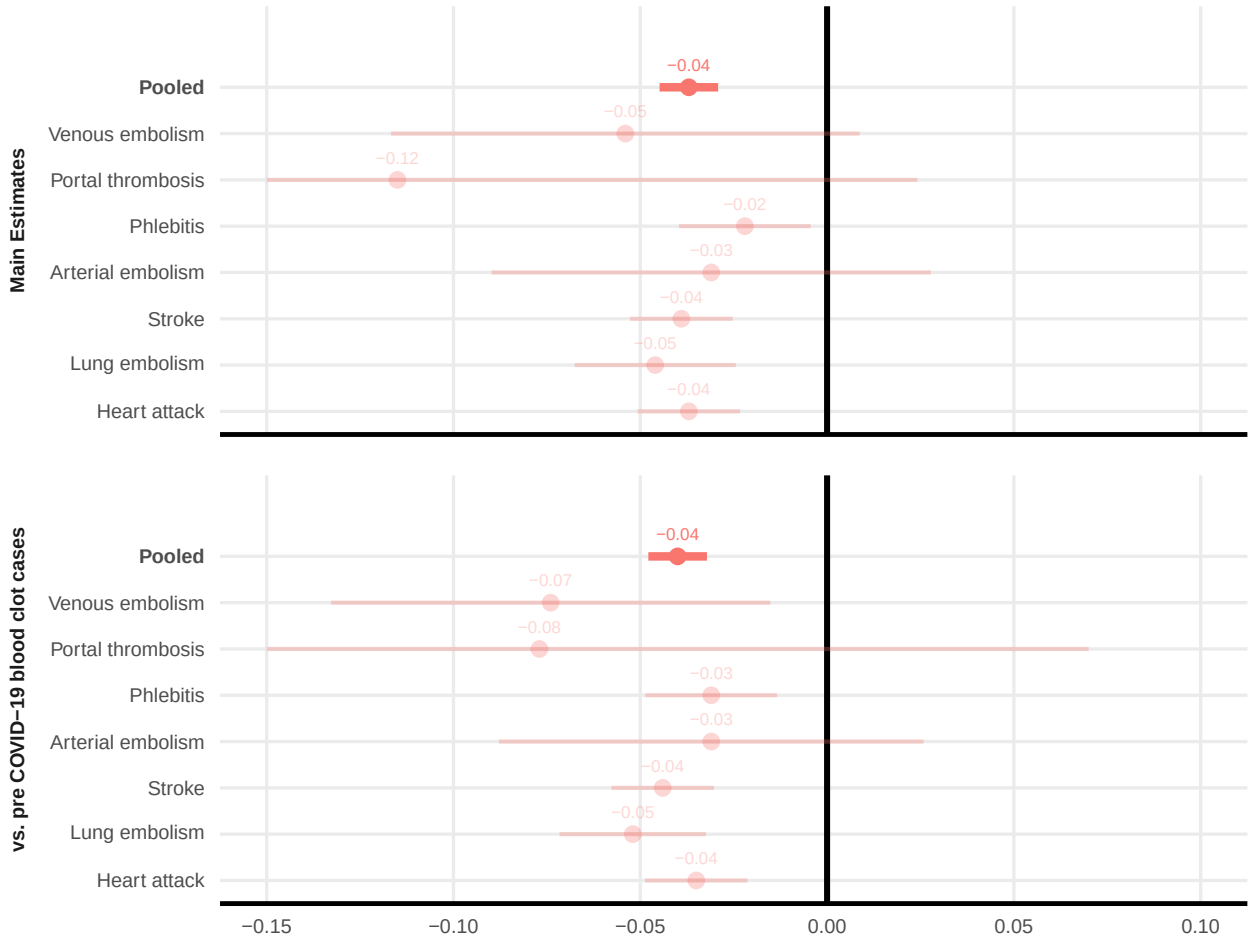

Figure B4: Vaccine Hesitancy and Media Exposure

Notes: This figure replicates the results in Figure 3, but contrasts coefficients from regressions run on matched samples where the control individuals are drawn from different candidate pools: (i) the full population (as in Figure 3), and (ii) individuals diagnosed with any blood clot diagnosis before the COVID-19 pandemic. Horizontal lines show 95% confidence intervals based on standard errors clustered at the match-pair level.

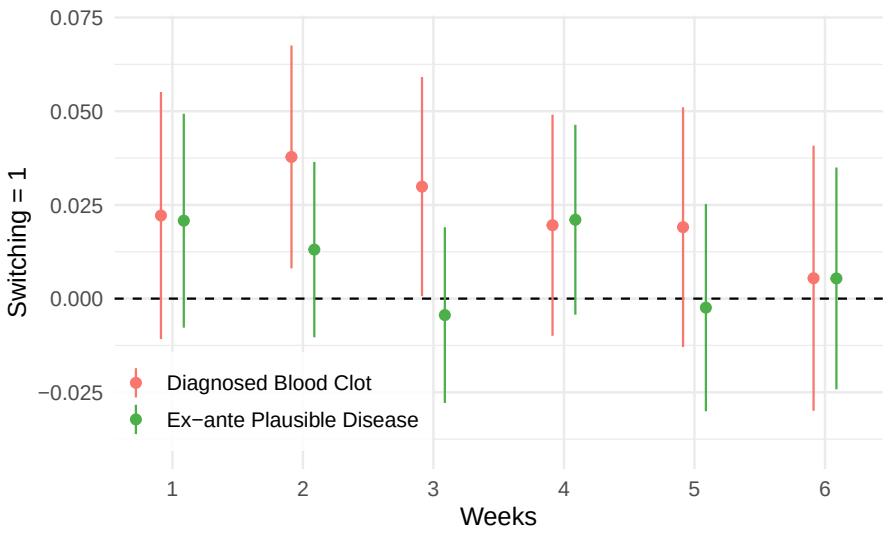

Figure B5: Heterogeneity by Time Elapsed Between Vaccination and blood Clot

Notes: This figure displays coefficient estimates of regressions run on a matched sample where the dependent variable is equal to 1 if an individual switches vaccine brand. The sample is split by the number of weeks elapsed between the last COVID-19 vaccine dose and developing a condition. Error bars represent 95% confidence intervals based on standard errors clustered at the match-pair level.

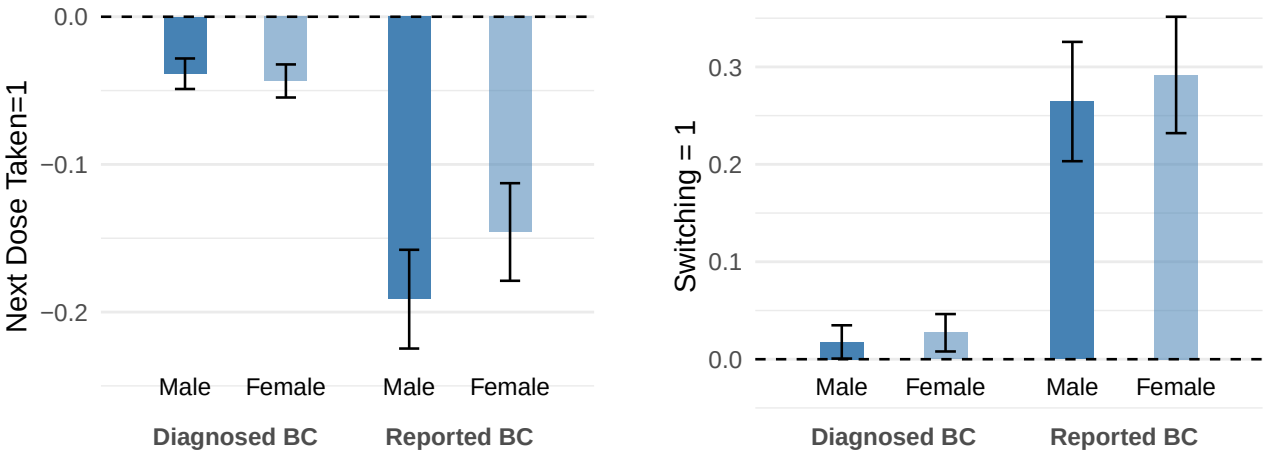

(a) Continuing Vaccination

(b) Switching Vaccines

Figure B6: Role of doctors' communication, heterogeneity by gender

Notes: This figure reports coefficient estimates from regressions estimated on the matched sample split up by gender. Diagnosed BC indicates individuals diagnosed with a blood clot within six weeks of receiving a COVID-19 vaccine (or their matched controls), while Reported BC further restricts the sample to cases reported as suspected adverse events by a healthcare professional. **Panel (a)** uses an indicator for continued COVID-19 vaccination as the dependent variable. **Panel (b)** uses an indicator for switching vaccine brand as the dependent variable. Error bars represent 95% confidence intervals based on standard errors clustered at the match-pair level.

Appendix C. Description of Matching

Propensity scores are computed using an XGBoost model with a binary logistic objective. Categorical variables are one-hot encoded, including missing-value indicators. The model is trained using 5-fold stratified cross-validation. To prevent overfitting, we employ early stopping within each fold of cross-validation after 10 rounds of no performance improvement on the test data set. Hyperparameters are set as follows: a learning rate (η) of 0.1, a maximum tree depth of 4, a minimum child weight of 1, a subsample fraction of 0.8, and a column subsampling fraction of 0.8. These settings provide moderate regularization and control overfitting while capturing nonlinear relationships.

We match each treated unit to one control unit. The large number of candidate matches for each individual in combination with the tight blocking conditions makes the global problem of finding close propensity score-neighbors without replacement challenging. Instead, we implement matching with replacement. Given the large number of candidate control individuals, this is unlikely to be a concern for inference.

The propensity scores are computed based on the following variables:

- **Gender** Indicator for male/female based on gender assigned at birth.
- **Birth year** Year of birth.
- **Origin** Defined as Sweden, EU, or rest of the world based on the country of birth of the individual and their parents.
- **Years of schooling** Proxy for years of schooling based on highest degree attained.
- **Field of education.** Two-digit education fields from Svensk utbildningsklassifikation (SUN), resulting in 25 categories.
- **Days unemployed** Average number of days unemployed per year, computed separately for 2006–2010, 2011–2015, 2016–2020.
- **Income** Average taxable income per year, computed separately for 2006–2010, 2011–2015, 2016–2020.
- **Doctor in family** Indicator equal to 1 if the individual or a family member holds a medical (physician) or nursing degree, based on SUN.
- **Municipality** Municipality of residence (290 categories).
- **Healthcare visits** Average number of healthcare visits per year in 2015–2020.
- **Prescription drugs** Average number of prescription drugs dispensed per year in 2015–2020.
- **Swine flu vaccine uptake** Indicator equal to 1 if the individual received the swine flu vaccine during the 2009–2010 pandemic (observed for 24% of the main sample).
- **COVID-19 risk group** Indicator equal to 1 if the individual belonged to a COVID-19 risk group and was prioritized for vaccination, based on prior diagnoses.

Appendix D. Data & Variable Definitions

D.1 Defining Other Acute Conditions

To identify Other Acute Conditions, we begin by selecting diagnoses that are sufficiently common to generate meaningful statistical power—those with at least 1,000 cases per year on average between 2015 and 2022 (596 ICD-10 codes). We then restrict to diagnoses for which more than half of cases originate in specialized care (93 codes), ensuring that the conditions reflect medically significant events rather than routine primary-care contacts. Next, we exclude ICD chapters that are clearly unsuitable for our context: circulatory diseases targeted by specific treatments (chapter I), external causes of injury (chapters S–Y), non-diagnoses such as symptoms or health-status factors (chapters R, U, Z), pregnancy-related and congenital conditions (chapters O, P, Q), cancers (chapters C, D), and mental-health conditions (chapter F), leaving 31 codes. Finally, we remove a small set of individual diagnoses that are either dominated by COVID-specific patterns (e.g., certain respiratory infections), reflect chronic or lifestyle conditions (e.g., diabetes complications, COPD, nutritional deficiencies, alcohol-related liver disease), represent post-surgical complications, or are too closely related to blood clots to serve as a meaningful comparison. This results in a set of 19 ICD-10 codes that plausibly could be mistaken for vaccine-related by a non-expert but are not known to be mechanistically linked to COVID-19 vaccination.

Table D1: Distribution of diagnoses among individuals developing other acute conditions following vaccination

Diagnosis (ICD-10)	**Share (%)**	**Number of individuals**	**Mortality rate (% deceased before 2023)**
Infectious agents (B95–B98)	27.6	9,359	24.2
Other bacterial diseases (A30–A49)	22.6	7,650	32.6
Appendicitis (K35–K38)	16.4	5,552	1.2
Acute kidney failure (N17–N19)	8.5	2,881	41.1
Bacterial pneumonia (J15)	7.1	2,415	36.0
Pleural effusion (J90)	6.0	2,040	47.0
Unspecified kidney failure (N19)	2.6	882	40.6
Peritonitis (K65–K67)	2.4	825	18.5
Respiratory failure (J96)	1.6	546	58.1
Pulmonary oedema (J81)	1.5	497	46.6
Aspiration pneumonitis (J69)	0.9	317	60.8
Pneumococcal pneumonia (J13)	0.7	248	18.4
Respiratory failure, unspecified (J96.9)	0.7	235	22.6
Hepatic failure (K72.9)	0.4	138	57.0
Pyothorax / empyema (J86)	0.3	101	22.6
Others	0.6	201	45.3
Total	100.0	33,887	26.3

Notes: This table displays the distribution of other acute conditions in the sample of individuals living in Sweden in 2021–2023 (columns 1 and 2). Diagnoses are included if recorded within 42 days of the first or second COVID-19 vaccination. Column 3 presents all-cause mortality for each diagnosis code, measured as the share deceased before 2023 within each category.

D.2 Defining blood clot-related side effects

The surveillance system contains brief descriptions of side effects following the MedDRA classification system. We use these descriptions to classify reports of suspected side effects as blood clot-related. In particular, we use the following terms:

- **Platelet disorders:** thrombocytopenia
- **Cardiac arrhythmias:** atrial fibrillation
- **Coronary artery disorders:** acute myocardial infarction; myocardial infarction
- **Retina, choroid and vitreous haemorrhages and vascular disorders:** retinal artery occlusion; retinal vascular thrombosis; retinal vein occlusion
- **Infections, pathogen unspecified:** sepsis
- **Central nervous system vascular disorders:** cerebral artery embolism; cerebral infarction; cerebral thrombosis; cerebrovascular accident; ischaemic stroke; cerebral venous sinus thrombosis; transient ischaemic attack
- **Pulmonary vascular disorders:** pulmonary embolism; pulmonary infarction
- **Embolism and thrombosis:** embolism; thrombosis; venous thrombosis; deep vein thrombosis; peripheral embolism; superficial vein thrombosis; thrombophlebitis; venous thrombosis limb

Out of these reported symptoms, an overwhelming majority of symptoms in our regression sample are for pulmonary embolism and deep vein thrombosis.

D.3 Defining date of availability

We adopt a data-driven approach to determine when the vaccine first becomes available to an individual, defining availability dates separately for each birth-year–by–healthcare-region combination. For each birth-year × region cell, let $f(i)$ denote the vaccination date of the i-th individual to be vaccinated (excluding healthcare workers). We define the date of availability as $f(i^*)$, where

$$i^* = \arg \min_{i \leq N-50} \{f(i+50) - f(i)\}.$$

That is, the availability date is the vaccination date at which the time span between individual i and the individual vaccinated fifty places later is minimized.