



A game-theoretic analysis of childhood vaccination behavior: Nash versus Kant

Philippe de Donder, Humberto Llavador, Stefan Penczynski, John E Roemer,
Roberto Vélez

► To cite this version:

Philippe de Donder, Humberto Llavador, Stefan Penczynski, John E Roemer, Roberto Vélez. A game-theoretic analysis of childhood vaccination behavior: Nash versus Kant. 2021. hal-03504644

HAL Id: hal-03504644

<https://hal.science/hal-03504644>

Preprint submitted on 29 Dec 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

December 2021

“A game-theoretic analysis of childhood vaccination
behavior: Nash versus Kant”

Philippe De Donder, Humberto Llavador, Stefan Penczynski,
John E. Roemer and Roberto Vélez

“A game-theoretic analysis of childhood vaccination behavior: Nash versus Kant”^{*,*}

by

Philippe De Donder, Toulouse School of Economics (CNRS)[†]

Humberto Llavador, Universitat Pompeu Fabra and Barcelona School of Economics[‡]

Stefan Penczynski, University of East Anglia[§]

John E. Roemer, Yale University^{**,}

Roberto Vélez, Centro Estudios Espinosa Yglesias, Mexico City^{††}

December 2021

Supporting Information

Additional Supporting Information can be found in the Online Appendix at <https://bit.ly/3xrHJ3R>

* We thank Peter Aronow, Nicholas Christakis, Mikhael Guy, Markus Jantti, Walter Orenstein, Nicolas Pistolesi, François Poinas, Kenneth Scheve, Joaquim Silvestre, Milan Svolik, Roberto Veneziani, Rodolfo de la Torre and seminar participants at the Aix-Marseille School of Economics, Paris School of Economics and Brown University for their advice and comments. Collin Schumock, Austin Jan, and Joon Lee, all from Yale University, provided excellent research assistantship.

[†] Philippe De Donder acknowledges the French ANR under grant ANR-17-EURE-0010 (Investissements d'Avenir program) for financial support. philippe.dedonder@tse-fr.eu

[‡] Humberto Llavador acknowledges financial support from the Spanish Agencia Estatal de Investigación (AEI), through the Severo Ochoa Programme for Centres of Excellence in R&D (Barcelona School of Economics CEX2019-000915-S) and grants AEI/FEDER, UE-PGC2018-096370-B-I00 and PID2020-115044GB-I00//AEI/10.13039/501100011033. humberto.llavador@upf.edu

[§] s.penczynski@uea.ac.uk

^{**} John Roemer acknowledges the Institute for Social and Policy Studies at Yale University for financial support. john.roemer@yale.edu

^{††} rvelezg@ceey.org.mx

Abstract

Whether or not to vaccinate one's child is a decision that a parent may approach in several ways. The vaccination game, in which parents must choose whether to vaccinate a child against a disease, is one with positive externalities (herd immunity). In some societies, not vaccinating is an increasingly prevalent behavior, due to deleterious side effects that parents believe may accompany vaccination.

The standard game-theoretic approach assumes that parents make decisions according to the Nash behavioral protocol, which is individualistic and non-cooperative. Because of the positive externality that each child's vaccination generates for others, the Nash equilibrium suffers from a free-rider problem. However, in more solidaristic societies, parents may behave cooperatively –they may optimize according to the Kantian protocol, in which the equilibrium is efficient. We test, on a sample of six countries, whether childhood vaccination behavior conforms better to the individualistic or cooperative protocol. In order to do so, we conduct surveys of parents in these countries, to ascertain the distribution of beliefs concerning the subjective probability and severity of deleterious side effects of vaccination.

We show that in all the countries of our sample the Kant model dominates the Nash model. We conjecture that, due to the free-rider problem inherent in the Nash equilibrium, a social norm has evolved, quite generally, inducing parents to vaccinate with higher probability than they would in the non-cooperative solution. Kantian equilibrium offers one precise version of such a social norm.

Keywords: Kantian equilibrium, Nash equilibrium, vaccination, social norm.

JEL code: C72 D62 D63 I12

1 Introduction

Vaccination against childhood diseases has improved child health and life expectancy dramatically over the last fifty years. Researchers from the US Center for Disease Control and Prevention (CDC) and the World Health Organization (WHO) write that in 2017, 110,000 children died of measles infection globally, and that in the period 2000-2017, 21 million lives were saved by measles vaccination (Dabbagh, Laws et alii, 2018, Table 2). The fraction of children globally who are vaccinated against measles rose in this period from 72% to 85%. Sweden and China have vaccination coverage rates in 2019 of 97% and 99%, respectively (World Bank, <https://data.worldbank.org/indicator/SH.IMM.MEAS?view=map>).

Our interest in this article, however, is not epidemiological, but rather behavioral. Vaccination is a choice in which cooperation among the population is important. Because of the phenomenon of herd immunity, the choice of each person to vaccinate has a positive welfare consequence for others. As the coverage rate of vaccination for a disease increases, it becomes more difficult for the virus to find hosts, and hence the probability that unvaccinated persons contract the disease decreases. Here, we model vaccination behavior as a game, in which the strategy of parents is to choose whether or not to vaccinate their child, or, in a more general version, a parent's mixed strategy is a probability that she will vaccinate her child.

The main equilibrium concept that economic theory employs to analyze this behavior is Nash equilibrium. A Nash equilibrium is a profile of vaccination probabilities for the population of parents such that, given the coverage rate, no parent can increase her expected welfare by altering her strategy (vaccination probability). Computing expected welfare requires trading off one possible lottery for another. If the parent does not vaccinate her child, the child may contract the disease, but is less likely to do so if the coverage rate is high; on the other hand, if vaccinated the child may be healthy, or suffer a possibly serious side effect of vaccination, or so the parent may believe. Expected welfare of the parent is the probability of vaccinating multiplied by the expected utility of the 'side-effect lottery,' plus the probability of not vaccinating times the expected utility of the 'unvaccinated' lottery.

The vaccination game, that we model explicitly below, has the property of being a *monotone increasing game*, which means that as long as the parent's probability of vaccinating her child is less than one, her expected utility will be greater the higher is the coverage rate (the more that other parents vaccinate). The general definition of a monotone increasing game is one in which each player's payoff is an increasing function of the strategies of the other players. The higher the probability that others vaccinate, the higher is my welfare, due to the increase in herd immunity. Now a Nash equilibrium of a monotone increasing game always suffers from the *free-rider problem*. People do not vaccinate 'enough' in the Nash equilibrium. This is because, in that equilibrium, no parent takes into account the positive effect of *her* choice to vaccinate on the welfare of other parents (or children). The consequence is *Pareto inefficiency of the equilibrium*: all could increase their welfare if all increased their probabilities of vaccination from the Nash equilibrium strategies. This is a generic welfare pathology of Nash equilibrium in monotone increasing games.

Nash equilibrium models the ethos of 'going it alone;' each parent takes as given – as part of her environment – the strategies of other parents, and asks only if she should alter her own strategy. She treats other parents' behaviors as fixed, and it is this assumption that gives rise to the free-rider problem.

There is, however, another equilibrium concept that models the ethos of *cooperation*, in contrast to the ethos of 'going it alone.' This is called Kantian equilibrium, which comes in different varieties. The simplest one applies only when all players in the game are identical. In our case, this would mean that all players have the same beliefs about the relevant probabilities of death from measles and of possibly suffering a bad side effect from the vaccination. In the case of identical players, each player asks herself: What is the action I would like everyone to take? Would I prefer everyone to vaccinate, or nobody to vaccinate, or everyone to vaccinate with the same probability a ? If all parents were identical, they will all agree on the answer to this question: it will be, for some number $a^* \in [0,1]$, that all players vaccinate with probability a^* . The consequence will be that (in a large population) the coverage rate will be a^* . The important result is that this '*Kantian*' equilibrium is always Pareto efficient in a monotone increasing game: the free-rider problem does not occur.

It's important to note that the Kantian optimization protocol does not rely on altruism of parents. When a parent asks 'what strategy should we *all* play?' she is forced to take into account the external effect on her own welfare implied by the actions of others. If we are all to play the same strategy, and if that is a low probability of vaccinating, the coverage rate will clearly be lower than otherwise, which is bad for *my* child. The question the player poses induces her to take into account the positive externality of vaccination. The consequence, though not obvious, is that Pareto efficiency is achieved in the equilibrium.

The 'simple Kantian equilibrium' just described will typically not exist if players are different –if they have different beliefs about the trade-offs required to evaluate whether or not to vaccinate. In this case, Kantian equilibrium takes a different form. Suppose there is a given profile of probabilities of vaccination among a population of parents. Now each parent asks: Would I like to *re-scale the entire vector of probabilities* by *some* positive (or zero) multiplicative constant? If all answer 'no' to this question, the profile is a *multiplicative Kantian equilibrium*.

Notice the difference between the question the Kantian player asks and the Nash player asks. The Kantian asks if she would like to re-scale *everyone's* vaccination probability, while the Nash parent asks if she would like to change *only her own* probability. The Kantian question builds symmetry into the equilibrium concept that is absent in Nash equilibrium. It turns out that, if all probabilities in the equilibrium profile are positive, then the Kantian equilibrium is Pareto efficient. Again, the free-rider problem dissolves.¹

This equilibrium concept is called Kantian because it models the idea to 'take an action if and only if you would will that it be universalized,' the categorical imperative of Immanuel Kant. Mathematically, the concept is a first cousin of Nash equilibrium, although the consequences are quite different in the two cases. Nash's approach, of defining an equilibrium of a game when no player wishes to choose another strategy

¹ The theorems concerning the Pareto efficiency of Kantian equilibrium are presented in Roemer (2019, chapters 2 and 3).

profile from a set of counterfactual profiles, can be altered to model not competition among players, but cooperation among them.

We can now state our project in this article. We will calibrate the parameters of the vaccination game using surveys that interview parents in a set of countries about their beliefs and their vaccination choices. We will then compute the Nash equilibrium and the Kantian equilibrium of the game in each country. This will consist of two profiles of vaccination probabilities in the country, and their implied equilibrium coverage rates. We will then ask which of these equilibria appears to better explain observed vaccination behavior in the country. Do parents appear to be ‘going it alone’ or cooperating? Of course, the reality is surely that some people go it alone and some behave cooperatively, but we will not attempt to analyze a model that is so nuanced: we will be satisfied with the simpler question just posed.

Section 2 presents a random utility model of vaccination and the equilibrium theory. Section 3 describes the data. Section 4 presents our method of estimation, that is, of testing whether the Nash or Kantian model better explains vaccination behavior in six countries. Section 5 presents our major finding: in all six countries, the Kantian model performs significantly better than the Nash model in explaining behavior. Section 6 argues that the reason this is true is that, due to the herd-immunity effect, a social norm has evolved to vaccinate one’s children, and this engenders vaccination rates that are uniformly greater than those predicted by Nash equilibrium. Section 7 offers a short conclusion. The Online Appendix presents details that are elided in the main text.

2 A random utility model of vaccination behavior

2.1 The set-up

We model the problem of the parent who must decide whether or not to vaccinate her child against measles. For now, we assume there are no laws or regulations mandating vaccination. If the child is not vaccinated, there is the possibility that he will contract measles and possibly die or suffer a debilitating illness. If he is vaccinated, he will either be healthy and protected from measles, or may suffer a side effect from the vaccination of some severity (or so the parent may believe).

We define three states of the child's health: healthy (H), suffering a possibly severe side effect from an inoculation (E), or contracting measles and possibly suffering a very severe outcome or death (D). Table 1 presents the probabilities of the three health states if vaccinated and if not vaccinated, and the von Neumann –Morgenstern utilities of the parent (the decision maker) based upon the child's health outcome.

Table 1. Utilities and probabilities of health states.

	Healthy	Side effect	Death/Severe Disability
Utility	1	u	0
Probability if not vax	$1 - p_0$	0	p_0
Probability if vax	$1 - p$	p	0

The utilities of the states Healthy and Death are a normalization that fixes the von Neumann-Morgenstern utility function of the parent. The utility u from the possible side effect is strictly between zero and one. We call the ordered pair (p, u) the parent's *type*; it is her beliefs about the utility-relevant facts concerning the side effect of vaccination. p_0 is probability of death or severe disability conditional upon contracting measles. We will take p_0 to be common knowledge of parents. We assume the population is characterized by a Beta distribution Q of (p, u) defined on the unit square. (Thus, we assume that $0 < u < 1$ for all types.) The parent's *mixed strategy* will be a number $a \in [0, 1]$, the probability with which she will vaccinate her child. The parent's (von Neumann Morgenstern) expected utility is defined on the ordered pair (a, \bar{a}) , where \bar{a} is the *coverage rate* in the population, defined as the average probability of vaccination across all (relevant) parents. There is a *probability function* $\pi: [0, 1] \rightarrow [0, 1]: \bar{a} \rightarrow \pi(\bar{a})$ which gives the probability that an unvaccinated child will contract measles if the coverage rate is \bar{a} . Herd immunity is modeled by supposing that π is a strictly decreasing, continuous function. Expected utility for a parent of type (p, u) is given by:

$$V_{(p,u)}(a, \bar{a}) = a \overbrace{\left((1-p) \cdot 1 + pu + \varepsilon \right)}^{\text{ex utility if vacc}} + (1-a) \overbrace{\left(\pi(\bar{a})(1-p_0) + (1-\pi(\bar{a})) \right)}^{\text{ex utility if not vacc}} \quad (2.1)$$

The number ε is the realization of a random variable distributed according to a distribution function $L(\cdot)$ on support \mathbb{R} , which is drawn i.i.d. across all parents. It is assumed that 99% of the support of L lies on the non-negative real numbers: this models a positive utility saltus that the parent receives if she vaccinates her child –because she is doing what physicians and society recommend, what most of her neighbors are doing, and so on. The main motivation for inserting this random element into utility is that it will guarantee that the Nash and Kant equilibrium strategies all lie in the open interval $(0,1)$, a property that is essential for our estimation strategy (see Section 4 below). Thus, the data of the problem are $\{Q, p_0, \pi(\cdot), L(\cdot)\}$.

We call the set of parents of a given type (p, u) a *tranche*. It is assumed, in particular, that ε is distributed i.i.d. within every tranche. This means that we will observe the behavior of a type (p, u) as a mixed strategy, even if every member of the tranche has a pure strategy, as long as the effect of the random variate differs across individuals. The statistician will only observe the average probability of vaccination within each tranche, which we will denote $a(p, u)$. This is to be thought of as the fraction of those of type (p, u) who decide to vaccinate, depending on their draw of the utility bump ε .

2.2 Nash equilibrium

A *Nash equilibrium of the game*, given a realization of the random variate L , is an action of ‘vaccinate’ or ‘do not vaccinate’ for every individual within every type given by:

$$\text{vaccinate} = \begin{cases} 1, & \text{if } \frac{dV_{(p,u)}}{da} \equiv p(u-1) + p_0\pi(\bar{a}^N) + \varepsilon > 0 \\ 0, & \text{if } \frac{dV_{(p,u)}}{da} \equiv p(u-1) + p_0\pi(\bar{a}^N) + \varepsilon < 0 \end{cases}, \quad (2.2)$$

Where \bar{a}^N is the fraction of individuals who vaccinate in equilibrium². Formally, we say that the Nash equilibrium is a strategy $\alpha^N(p, u, \varepsilon)$ for each individual (p, u, ε) and a coverage rate \bar{a}^N such that:

² We can ignore the null set of types for which $p(u-1) + p_0\pi(\bar{a}^N) + \varepsilon = 0$

$$\alpha^N(p, u, \varepsilon) = \begin{cases} 1 & \text{if } \varepsilon > p(1-u) - p_0\pi(\bar{a}^N) \\ 0 & \text{if } \varepsilon < p(1-u) - p_0\pi(\bar{a}^N) \end{cases} \quad (2.3)$$

and:

$$\bar{a}^N = \int_{(p,u)} \int_{p(1-u)-p_0\pi(\bar{a}^N)}^{\infty} dL(\varepsilon) dQ(p, u) = \int \left(1 - L(p(1-u) - p_0\pi(\bar{a}^N))\right) dQ(p, u) \quad (2.4)$$

□ Thus, vaccinate if and only if $\varepsilon > p(1-u) - p_0\pi(\bar{a}^N)$, an event that occurs (in the (p, u) tranche) with probability $1 - L(p(1-u) - p_0\pi(\bar{a}^N))$. The fraction of this tranche that vaccinates is:

$$a^N(p, u) = \int_{p(1-u)-p_0\pi(\bar{a}^N)}^{\infty} \alpha^N(p, u, \varepsilon) dL(\varepsilon) = 1 - L(p(1-u) - p_0\pi(\bar{a}^N)) \quad (2.5)$$

Equation (2.4) is a single equation in the unknown \bar{a}^N . We solve it for \bar{a}^N , and then compute the Nash equilibrium strategy profile from equation (2.5). Note that it appears as if the type (p, u) has a (single) mixed strategy, $a^N(p, u)$.

It is also noteworthy that, because the support of L is the entire real line, every Nash and Kantian strategy is in the open interval $(0,1)$. This will be an important fact in what follows.

2.3 Kantian equilibrium

It will be convenient to define:

$$p^*(u, \bar{a}) = \frac{p_0 p(\bar{a})}{1-u}. \quad (2.6)$$

As discussed in Section 1, a multiplicative Kantian equilibrium of a game in normal form among M players with payoff functions $\{V^i\}$ is a strategy profile $a(p, u)$ and a coverage rate $\bar{a} = \int a(p, u) dQ(p, u)$ such that no player would like to re-scale the profile by any non-negative factor, that is:

$$[\forall(p, u)] \left[1 = \operatorname{argmax}_{\rho \geq 0} V_{(p,u)}(\rho a(p, u), \rho \bar{a}) \right]. \quad (2.7)$$

A profile of vaccination strategies $\alpha^K(p, u, \varepsilon)$, where $\bar{a}^K = \int \alpha^K(p, u, \varepsilon) dL(\varepsilon) dQ(p, u)$ is a *multiplicative Kantian equilibrium of the vaccination game after the random variate L is realized* if no player (p, u, ε) would like to re-scale the entire profile by any non-negative factor ρ (see equation (2.9)).

What is the domain of ρ for a player (p, u, ε) –that is, in the game where the draw from the random variate L has been realized? At a Kantian equilibrium, a player (p, u, ε) must choose $\rho \in [0, 1/(\alpha^K(p, u, \varepsilon))]$ –that is, she cannot contemplate a re-scaling of the profile that gives her a vaccination probability greater than one. Furthermore, she must truncate at one, if need be, the re-scaling for other players so that they are not required to play infeasible strategies.

We must distinguish carefully between the equilibrium after the random variable L has assigned a value ε to every player, and what the statistician observes, not knowing the realization of L . Since at the observed equilibrium there will be players with all values of ε at a given (p, u) in the support of Q , and these players will have different strategies $\alpha^K(p, u, \varepsilon)$, what the statistician will observe is that the (p, u) tranche is playing a mixed strategy:

$$a^K(p, u) = \int \alpha^K(p, u, \varepsilon) dL(\varepsilon) . \quad (2.8)$$

Note that $\bar{a}^K = \int a^K(p, u) dQ(p, u) = \bar{\alpha}^K$, because we have already integrated over both (p, u) and ε in the definition of $\bar{\alpha}^K$. \bar{a}^K or $\bar{\alpha}^K$ is the vaccination rate in the population, observed by the statistician.

If a player/parent contemplates a re-scaling of a strategy profile $\alpha(p, u, \varepsilon)$ by a factor ρ that is larger than one but very close to one, the fraction of players who vaccinate (in the rescaled profile) will be very close to $\rho \bar{a} + (1 - \rho)Q(A)$, where A is the set of voters (p, u, ε) for whom $\alpha(p, u, \varepsilon) = 1$. This is the result of truncating the rescaled policies so that no player is required to vote with a probability greater than unity.

Thus, the expected utility of the parent in a rescaled profile is, for ρ close to one, given by:

$$\begin{aligned} \tilde{V}_{(p, u, \varepsilon)}(\alpha, \bar{a}, \varepsilon; \rho) = & \\ & \left\{ \begin{array}{ll} \overbrace{\rho \alpha \left((1-p) \cdot 1 + pu + \varepsilon \right) + (1-\rho \alpha) \left(\pi(\rho \bar{a} + (1-\rho)Q(A))(1-p_0) + \left(1 - \pi(\rho \bar{a} + (1-\rho)Q(A)) \right) \right)}^{\text{ex utility if not vacc}} & \text{if } \rho > 1 \\ \overbrace{\rho \alpha \left((1-p) \cdot 1 + pu + \varepsilon \right) + (1-\rho \alpha) \left(\pi(\rho \bar{a})(1-p_0) + \left(1 - \pi(\rho \bar{a}) \right) \right)}^{\text{ex utility if not vacc}} & \text{if } \rho \leq 1 \end{array} \right. \end{aligned} \quad (2.9)$$

By definition, a profile $\alpha^K(p, u, \varepsilon)$ is a Kantian equilibrium of the game where L has been realized exactly when:

$$\text{for all } (p, u, \varepsilon), 1 = \underset{0 \leq \rho \leq \left(\frac{1}{\alpha^K(p, u, \varepsilon)}\right)}{\operatorname{argmax}} \tilde{V}_{(p, u, \varepsilon)}(\alpha^K, \bar{a}^K, \varepsilon; \rho) \quad (2.10)$$

What are the first-order conditions associated with the optimization (2.10)? Note that $\tilde{V}_{(p, u, \varepsilon)}$ is not differentiable at $\rho = 1$, where there is a kink in the function. So, the f.o.c.s guaranteeing (2.10) if $\alpha^K(p, u, \varepsilon) < 1$ are:

$$(\forall \alpha) \quad \left. \frac{d^-}{d\rho} \right|_{\rho \nearrow 1} \tilde{V}_{(p, u, \varepsilon)} \geq 0 \text{ and } (\forall \alpha < 1) \quad \left. \frac{d^+}{d\rho} \right|_{\rho \searrow 1} \tilde{V}_{(p, u, \varepsilon)} \leq 0, \quad (2.11)$$

where $\frac{d^-}{d\rho}$ and $\frac{d^+}{d\rho}$ are the left-hand and right-hand derivatives, respectively. If $\alpha^K(p, u, \varepsilon) = 1$, the f.o.c.

is $\left. \frac{d^-}{d\rho} \right|_{\rho \nearrow 1} \tilde{V}_{(p, u, \varepsilon)} \geq 0$ alone, because the relevant domain of ρ in the second expression in (2.11) becomes $[0, 1]$.

Now we compute:

$$\begin{aligned} \left. \frac{d^-}{d\rho} \right|_{\rho \nearrow 1} \tilde{V}_{(p, u, \varepsilon)} &= \alpha(p(u-1) + \varepsilon + p_0\pi(\bar{a}) + p_0\pi'(\bar{a})\bar{a}) - p_0\pi'(\bar{a})\bar{a} \geq 0 \\ \left. \frac{d^+}{d\rho} \right|_{\rho \searrow 1} \tilde{V}_{(p, u, \varepsilon)} &= \alpha(p(u-1) + \varepsilon + p_0\pi(\bar{a}) + p_0\pi'(\bar{a})(\bar{a} - Q(A))) - p_0\pi'(\bar{a})(\bar{a} - Q(A)) \leq 0 \end{aligned} \quad (2.12)$$

Both conditions in (2.12) must hold if $\alpha^K(p, u, \varepsilon) < 1$. If $\alpha^K(p, u, \varepsilon) = 1$, the first condition only need hold, and it becomes:

$$\varepsilon \geq (1 - u)(p - p^*(u, \bar{a})).$$

Define $\alpha^K(p, u, \varepsilon)$ so that $\left. \frac{d^-}{d\rho} \right|_{\rho \nearrow 1} \tilde{V}_{(p, u, \varepsilon)} = 0$ when $\varepsilon < (1 - u)(p - p^*(u, \bar{a}))$, and define $\alpha^K = 1$

otherwise. From (2.12) this requires:

$$\alpha^K(p, u, \varepsilon) = \begin{cases} \frac{-p_0\pi'(\bar{a}^K)\bar{a}^K}{(1 - u)(p - p^*(u, \bar{a}^K)) - \varepsilon - p_0\pi'(\bar{a}^K)\bar{a}^K}, & \text{if } \varepsilon < (1 - u)(p - p^*(u, \bar{a})) \\ 1 & \text{if } \varepsilon \geq (1 - u)(p - p^*(u, \bar{a})) \end{cases} \quad (2.13)$$

The statistician sees only the average coverage rate for each tranche (p, u) . This is given by:

$$a^K(p, u) = \int_{-\infty}^{(1-u)(p-p^*(u, \bar{a}^K))} \frac{-p_0 \pi'(\bar{a}^K) \bar{a}^K}{(1-u)(p-p^*(u, \bar{a}^K)) - \varepsilon - p_0 \pi'(\bar{a}^K) \bar{a}^K} dL(\varepsilon) + 1 - L\left((1-u)(p-p^*(u, \bar{a}^K))\right). \quad (2.14)$$

Integrating over all (p, u) :

$$\bar{a}^K = \int \int_{-\infty}^{(1-u)(p-p^*(u, \bar{a}^K))} \frac{-p_0 \pi'(\bar{a}^K) \bar{a}^K}{(1-u)(p-p^*(u, \bar{a}^K)) - \varepsilon - p_0 \pi'(\bar{a}^K) \bar{a}^K} dL(\varepsilon) dQ(p, u) + \int \left[1 - L\left((1-u)(p-p^*(u, \bar{a}^K))\right)\right] dQ(p, u), \quad (2.15)$$

which is an equation in the single unknown \bar{a}^K . We solve it for the Kantian coverage rate, and then compute the equilibrium strategy profile from (2.14).

To check that the strategy profile α^K , so defined, is a Kantian equilibrium of the game that takes place after the random variable L has been realized, we need only check that (2.11) holds. By definition, we have

$\frac{d^-}{d\rho} \Big|_{\rho \nearrow 1} \tilde{V}_{(p, u, \varepsilon)} = 0$ when $\alpha^K < 1$. Now note from the formulae in (2.12) that:

$$\frac{d^+}{d\rho} \Big|_{\rho \searrow 1} \tilde{V}_{(p, u, \varepsilon)} = \frac{d^-}{d\rho} \Big|_{\rho \nearrow 1} \tilde{V}_{(p, u, \varepsilon)} + (1 - \alpha) p_0 \pi'(\bar{a}) Q(A). \quad (2.16)$$

But the term $(1 - \alpha) p_0 \pi'(\bar{a}) Q(A)$ is negative or zero, which implies that $\frac{d^+}{d\rho} \Big|_{\rho \searrow 1} \tilde{V}_{(p, u, \varepsilon)} \leq 0$, and so (2.11)

follows. Finally, note that when $\alpha^K = 1$, $\frac{d\tilde{V}}{d\rho} \Big|_{\rho \nearrow 1} > 0$, and the second condition in (2.12) becomes vacuous,

because the domain of ρ becomes $\{1\}$.

2.4 Comparison of Kantian and Nash vaccination equilibria

We noted in Section 1 that the vaccination game is a monotone increasing game. (Just check in equation (2.1) that $V_{(p, u)}$ is an increasing function of \bar{a} .) This is the mathematical consequence of herd immunity. It follows that the Nash equilibrium of the game will suffer from the free-rider problem, but

the multiplicative Kantian equilibrium will be Pareto efficient. Intuitively, people will vaccinate ‘too little’ in the Nash equilibrium. The precise consequence is this:

Theorem 1 $\bar{a}^K > \bar{a}^N$.

The coverage rate is greater in Kantian equilibrium.

Proof:

Suppose to the contrary that $\bar{a}^N \geq \bar{a}^K$. Then $p^*(u, \bar{a}^K) \geq p^*(u, \bar{a}^N)$ and this implies that the second term in the r.h.s. of equation (2.15) is greater than the r.h.s. of equation (2.5). A fortiori, $\bar{a}^K > \bar{a}^N$ because the first term on the r.h.s. of equation (2.15) is positive. This contradiction proves the claim. \square

In fact, we can say more. Note that although we have defined a parental type as an ordered pair of traits/beliefs (p, u) , in fact the population profile of traits can be more parsimoniously written as depending only on the single variable $w = p(1 - u)$. For we can write the Nash and Kantian equilibrium policies, from equations (2.5) and (2.14) respectively as:

$$\tilde{a}^N(w) = 1 - L(w - p_0\pi(\bar{a}^N)) \quad (2.17)$$

and:

$$\tilde{a}^K(w) = \int_{-\infty}^{w - p_0\pi(\bar{a}^K)} \frac{-p_0\pi'(\bar{a}^K)\bar{a}^K}{w - p_0\pi(\bar{a}^K) - \varepsilon - p_0\pi'(\bar{a}^K)\bar{a}^K} dL(\varepsilon) + 1 - L(w - p_0\pi(\bar{a}^K)) \quad .^3 \quad (2.18)$$

Since the domain of (p, u) is the unit square, the domain of w is $[0, 1]$. We can plot the difference of the two equilibrium profiles

$$\Delta\tilde{a}(w) = \tilde{a}^K(w) - \tilde{a}^N(w). \quad (2.19)$$

³ Note we have written the coverage rates in equations (2.17) and (2.18) as \bar{a}^N and \bar{a}^K . We could have written these as \bar{a}^N and \bar{a}^K . The coverage rates will be the same whether we integrate $dQ(p, u)$ or $d\tilde{Q}(w)$, where $\tilde{Q}(\cdot)$ is the distribution function of w induced by Q .

See Figure 1a in Section 5 below. When w is small then either p is small or u is close to one, or both, so the parent either believes that the probability of a severe side effect from vaccination is small, or if the side effect occurs, it is not severe (u close to one means the health status of a child with the side effect is close to full health). So parents with w close to zero will be likely to vaccinate according to our model and parents with w close to one will be likely not to vaccinate.

3 Producing the data

The data we require to compute the Kantian and Nash equilibria for a society are p_0 , the distribution Q of (p, u) , and the function $\pi(\cdot)$. We describe the choice of the logistic variate L below. We have administered the survey to adults aged 20 to 45 in the US, the UK, Germany, France, Canada, and Mexico. The survey is presented in Section III of the Online Appendix.

Table 2 shows some descriptive statistics.

Table 2. Descriptive statistics of the country surveys

	Canada	France	Germany	UK	US	Mexico
Mean age	35.3	34.7	33.1	33.9	32.2	31.3
Female %	49.1	53.5	54.3	52.3	60.3	55.6
Parents since 2011 %	32.3	56.1*	33.8	41	32.8	56.5
Measles vaccine %	88.9	90.3	89.3	89.2	82.9	96.8
Covid vaccine %	73.2	50	64.1	77.3	54.1	88.7
N	1052	1188	1146	1054	1210	1063

* In the French survey the question asked was "Do you have a child born in or before 2018?".

The probabilities p and p_0 representing the individual's beliefs are ascertained in a standard way in the questionnaire.

We estimate u by presenting the respondent with a series of binary choices over pairs of lotteries. This technique allows us to place the respondent's value of u in a relatively small interval within $[0,1]$. The method assumes the individual is an expected utility maximizer.⁴ We pose the question:

⁴ See Holt and Laury (2002) for a description of this approach.

- In the following scenario, would you prefer event A or event B:
 - A. For your child to have a bad side effect from a measles vaccination, or
 - B. For your child to face an unrelated risk in which he/she has a 99% chance of being healthy, and a 1% chance of dying.

Suppose the respondent answers B. If utility is normalized as in Table 1, then we conclude that $(0.99 \times 1 + 0.01 \times 0) = 0.99 > u$. Next, we ask:

- In the following scenario, would you prefer event A or event B:
 - A. For your child to have a bad side effect from a measles vaccination, or
 - B. For your child to face an unrelated risk in which he/she has a 95% chance of being healthy and a 5% chance of dying.

Suppose the respondent answers A. Then we conclude that $u > (0.95 \times 1 + 0.05 \times 0) = 0.95$, and hence we know that $u \in (0.95, 0.99)$. We assign this respondent a value of u chosen randomly from this interval.

Thus, we ascertain the respondent's value of u by posing a series of such questions about lottery choice.

We then fit a bivariate Beta distribution defined on $[0,1]^2$ to the respondents' values of (p, u) . The Beta distribution is calculated knowing the observed means and variances of p and u , and their covariance.

Table 3 presents these data for our six countries.

Table 3. Data from the country surveys

	Mean p	Var p	Mean u	Var u	Cov(p, u)	Median p_0
US	0.048	0.026	0.851	0.061	-0.009	0.003
UK	0.020	0.009	0.891	0.043	-0.004	0.002
Germany	0.020	0.010	0.863	0.054	-0.004	0.002
France	0.022	0.011	0.743	0.094	-0.002	0.003
Canada	0.017	0.052	0.874	0.052	-0.0009	0.001
Mexico	0.035	0.015	0.774	0.068	-0.0005	0.005

We choose the common value of p_0 from the survey to be the median response to the appropriate question on the survey. The median is a better choice than the mean value, as the latter is distorted by several very high and unreasonable values for p_0 .⁵

We comment on the value of p_0 , the median value of respondents' opinions on the probability of dying from a measles infection. Dabbagh, Laws et alii (2018, Table 1) report that in 2017, for the continent of Europe the actual value is $p_0 = 0.004 = 0.4\%$, slightly larger than the median respondent's opinion. Unfortunately, this article does not present the value for the United States. But for Africa, the reported value of p_0 has a point estimate of 0.66, and the lower-bound estimate in the 95% confidence interval is 0.31. Measles can be a deadly disease if medical care is poor.

The most severe side effect of MMR (measles, mumps, rubella) vaccination is aseptic meningitis, which occurs in 1 in 10 million cases. The probabilities p that respondents to our survey give are greater than this by four orders of magnitude; however, from the values of u respondents provide, they are on average viewing side effects as not terribly severe (a value of $u = 0.85$ says that good health is 15% reduced by the side effect). The possibly bad outcomes of measles are considerably worse, and include, besides death, anaphylaxis, febrile seizures, thrombocytopenic purpura and encephalitis (see Strebel and Orenstein, 2019, which also gives the probabilities). The anti-vaccination movement is often motivated by fears that vaccination may cause autism, which were falsely aroused in a 1998 article published in *Lancet*, later retracted by *Lancet* in 2010.

We use the following parametric form for the probability function:

$$\pi(\bar{a}) = (1 - \bar{a})^\gamma, \quad (3.1)$$

where \bar{a} is the observed measles vaccination coverage rate for the country. We chose the parameterization (3.1) as possibly the simplest functional form that gives a decreasing function passing through the points (0,1) and (1,0). In the Appendix A, we describe the precise definition of the function π and how we estimate

⁵ However, we report the mean values of p and u because these are used to fit the Beta distribution Q to the data.

γ . For Canada and the United States, we estimate $\gamma = 3.1$. For the UK, France, Germany and Sweden, we estimate $\gamma = 1.995$.⁶ We split our set of countries in two because the number of cases of measles in the last five years in Europe has been an order of magnitude larger than in North America (Canada and the US), despite the higher coverage rates enjoyed by the European countries. In particular, the number of measles cases in Germany in 2015 was several orders of magnitude higher than in Canada and the US, and this is doubtless because 2015 was the year during which Prime Minister Angela Merkel admitted one million refugees from the Middle East. We presume the infection process therefore differs between recent European experience and the North American, justifying different values of γ in equation (3.1).⁷

At the WHO-reported⁸ coverage rate of 0.916 for the US, the probability that an unvaccinated child in a given cohort in the US contracts measles before the age of five, defined as the number of measles cases in her birth cohort divided by the number of unvaccinated children in her cohort, is 4.5×10^{-4} , or about one in 2000.

The last year measles was endemic in the United States was 2000.⁹ The aforementioned WHO data set reports that in 2019, measles was endemic in Germany and France. It is probably also endemic in Mexico, although the data are incomplete. We cannot use the SIR model to compute the probability of contracting measles because this model is not applicable to analyzing very small occurrences of the disease that are quickly stamped out¹⁰. In any case, the SIR model will not give us a probability as a function of the

⁶ We had planned to include Sweden in our sample of countries, and so included it in the estimation of γ . Unfortunately, doing so was eventually not possible. Estimating the European value of γ without Sweden gives a value of 2.007. Based on the small difference between this value and 1.995, we elected not to re-run all the equilibrium calculations for the UK, France and Germany with $\gamma = 2.007$, a costly procedure.

⁷ Without morbidity data for Mexico, we use the European value of $\gamma = 1.995$. Nevertheless, results do not change significantly when using the North American value.

⁸ Data source https://apps.who.int/immunization_monitoring/globalsummary/.

⁹ A contagious disease is endemic if an outbreak induces a sequence of contagion that does not terminate within a year.

¹⁰ A useful description of the SIR model is found in Avery, Bossert et al (2020).

coverage rate only: in that model, the probability that a susceptible individual contracts the disease is a function of two numbers –for instance, the fraction of susceptible (uninoculated) individuals (S), and the fraction of recovered individuals (R).

Our definition of the function p as the probability that a child who is unvaccinated contracts measles by the age of five is meant to model the relevant probability that a parent needs in order to decide whether or not to vaccinate her child.

Table 4. Coverage rates for measles, five-year average (WHO)

Country	US	UK	Germany	France	Canada	Mexico
Coverage	91.6%	92.2%	97.2%	90.2%	89.6%	86%

It is important to note that measles vaccination in our six countries is, or was until recently, *de facto* voluntary. In the United States, there is no federal law requiring children be vaccinated –such laws are left to the states. All 50 states require children be vaccinated against measles before attending childcare or public school; however, all states permit exemptions for medical, religious, or reasons of conscience, and the standards are not strict. In Canada, vaccination policies are taken at the provincial level. Only three provinces (Ontario, New Brunswick and Manitoba) have legislated requirements; however, exemptions are granted on medical or religious grounds, or simply out of conscience in these provinces. In the UK, childhood vaccination is not mandatory. In Germany, a federal law now requires measles vaccination, but only since March 1, 2020. In our German survey, we asked parents whether they vaccinated or did not vaccinate their child prior to that date. In France, vaccination was only recommended prior to 2018, and in the French questionnaire, we asked parents for the vaccination status of their child prior to 2018. Mexico has no law requiring vaccination.

4 Estimation procedure

We wish to decide whether the Nash model or the Kantian model provides a better explanation of observed vaccination behavior in a country. We have samples of roughly 1000 (N) respondents for each

country. Each respondent is characterized by a triple (p, u, v) where $(p, u) \in [0,1]^2$ is the vector of respondent traits and $v \in \{0,1\}$ indicates that the respondent did (1) or did not (0) vaccinate her child. We call $\mathbf{v}^{obs, s^0} = (v^1, \dots, v^N)$ the *observation* or *observed vaccination behavior* of the original sample. The superscript in s^0 refers to the original survey sample for the country.

There are three sources of randomness in our models. First, there is a logistic variate L , 99% of whose mass lies on the positive real line (more below). Each parent who chooses to vaccinate draws a realization of this variate i.i.d. across individuals, which is interpreted as a positive saltus in utility that the parent enjoys if she vaccinates her child. Secondly, since the equilibrium strategies observed by the statistician in both the Nash and Kantian model are mixed strategies, there is a random process which must determine whether a player with an equilibrium strategy $a \in (0,1)$ chooses $v = 0$ or 1. Third, there is a ‘trembling hand’ introduced below: with some probability q each player, when choosing the action v , misreads the coin flip that determines what her behavior should be. (These trembles will be i.i.d.) The purpose of the first and third sources of randomness is to make the models more realistic, so as to achieve a better fit to the observed vaccination behavior, and to guarantee that the Nash and Kant equilibrium strategies are all strictly mixed strategies (lie in the open interval $(0,1)$). The second source is due to the mixed-strategy character of the equilibria.

4.1 The logistic variate L

It is useful for computation to have the support of L be the entire real line: this guarantees that all equilibrium strategies, Nash and Kant, are in the open interval $(0,1)$. This motivates our choice of a logistic distribution. See equations (2.5) and (2.14), which guarantee that the probabilities of vaccination are never zero or one when L ’s support is \mathbb{R} . We shall determine L by a single parameter, its mean value μ . The logistic variate is in fact characterized by two parameters, denoted (μ, β) . Denote by $L^{(\mu, \beta)}$ the c.d.f. of the logistic with parameters (μ, β) . Given μ , we choose β so that:

$$L^{(\mu, \beta)}(0) = 0.01; \tag{4.1}$$

that is, 99% of L 's mass is on the positive real line. Hence L is chosen from a single parameter family, where the parameter is μ .

4.2 Nash and Kantian equilibria

We will perform the estimation procedure outlined in this section for a large number B of bootstrapped samples, obtained from the original survey sample s^0 by sampling from it with replacement. Let the size of the mother sample s^0 and of all the bootstraps for a particular country be N . Here we describe the estimation procedure using the mother sample s^0 ; the identical procedure will be carried out for every bootstrap sample s .

Given the sample s^0 , we fit a bivariate Beta distribution Q^0 to the observed distribution of (p, u) . μ is chosen to be a small positive number. For any choice of μ , the logistic distribution $L^{(\mu, \beta)}$ is determined, see (4.1). Given L and Q^0 we can compute the Nash and Kantian equilibria of the vaccination game as described in Section 2. The Nash equilibrium is a profile of strategies (probabilities of vaccinating) $a^N(p, u; s^0, \mu)$ and the Kantian equilibrium is a profile of strategies $a^K(p, u; s^0, \mu)$.

Given these two equilibria, we can compute *the log likelihood of the observed vaccination behavior* \mathbf{v}^{obs, s^0} . This is defined, for the Nash equilibrium, as:

$$\Psi^{Na}(s^0, \mu, \mathbf{v}^{obs, s^0}) = \sum_{\{(p, u)|v=1\}} \log a^N(p, u; s^0, \mu) + \sum_{\{(p, u)|v=0\}} \log(1 - a^N(p, u; s^0, \mu)), \quad (4.2)$$

and for the Kantian equilibrium as:

$$\Psi^{Ka}(s^0, \mu, \mathbf{v}^{obs, s^0}) = \sum_{\{(p, u)|v=1\}} \log a^K(p, u; s^0, \mu) + \sum_{\{(p, u)|v=0\}} \log(1 - a^K(p, u; s^0, \mu)), \quad (4.3)$$

where the original sample is the collection of triples $\{(p, u, v)\}$.

Since the strategies are all in the open interval $(0, 1)$, the two log likelihood functions are well-defined. Because of precision problems in computation, we in fact encounter some zero values in the computation of $a^N(p, u)$. Rather than eliminating these respondents from the sample, we replace the zero values of $a^N(p, u)$ with $ra^K(p, u)$ where $r = \frac{\text{mean}_{\{(p, u)|a^N(p, u) > 0\}} [a^N(p, u)/a^K(p, u)]}{1}$. It will turn out that $r < 1$, because $a^N(p, u) < a^K(p, u)$ for all (p, u) .

4.3 Analyzing the sample

Next, we ask: Could it be that \mathbf{v}^{obs, s^0} can be explained as an outcome of Nash behavior, but amended by a trembling hand that causes each respondent to choose the opposite behavior from what the Nash coin-flip produces? Let's say the tremble occurs i.i.d. for each respondent with probability q . In this case, an agent (p, u) chooses to vaccinate ($v = 1$) with probability :

$$a^{*N}(p, u) = (1 - q)a^N(p, u) + q(1 - a^N(p, u)). \quad (4.4)$$

Suppose we run a large number, Λ , of trials with this model, all with the sample s^0 . The only thing that differs across trials is the realization of the coin flips that implement the tremble: the expected value of the coin flip for an agent (p, u) is always given by $a^{*N}(p, u)$ in (4.4). Denote the index of the trial by l . Define:

$$\mathbf{1}_q^l = \{(p, u) | a^{*N}(p, u) \text{ coinflip } l \rightarrow 1\}, \mathbf{0}_q^l = \{(p, u) | a^{*N}(p, u) \text{ coinflip } l \rightarrow 0\}$$

We ask: What log likelihood would *this* observed vaccination outcome have if we mistakenly thought the true Nash model (absent the coin-flip) were the correct model? That likelihood is given by:

$$\Psi(q, l; s^0, \mu) = \sum_{(p, u) \in \mathbf{1}_q^l} \log a^N(p, u; s^0, \mu) + \sum_{(p, u) \in \mathbf{0}_q^l} \log(1 - a^N(p, u; s^0, \mu)). \quad (4.5)$$

We are taking the log likelihood of the observed behavior from the trembling-hand coin-flip experiment and evaluating it with respect to the *pure* Nash model, without the trembling hand.

Next, we want to compute the expected value of $\Psi(q, l)$ over $l = 1, 2, \dots, \Lambda$. We can write the expected log likelihood of the experiment as the number of trials Λ becomes large as:

$$M(q) \equiv \lim_{\Lambda \rightarrow \infty} \frac{1}{\Lambda} \sum_{l=1}^{\Lambda} \Psi(q, l) = \sum_{(p, u)} a^{*N}(p, u) \log a^N(p, u) + (1 - a^{*N}(p, u)) \log(1 - a^N(p, u)). \quad (4.6)$$

This is the key step. It's true because if we look at the sums in (4.5) over all l , by definition, a given (p, u) will lie in the set $\mathbf{1}_q^l$ for a fraction $a^{*N}(p, u)$ of the Λ trials, as Λ becomes large. And (p, u) will lie in $\mathbf{0}_q^l$ a fraction $1 - a^{*N}(p, u)$ of the time.

Our strategy is to ask how large a tremble is needed to produce the log likelihood $\Phi^{Na}(s, \mu, \mathbf{v}^{obs, s^0})$. Our claim is: the smaller the tremble needed to 'rationalize' the observed vaccination behavior, the better

explanation the model provides of observed behavior. In other words, since we view the trembling hand as a device for inserting randomness into the Nash (or Kant) model, then the less randomness required to explain the observed behavior, the better the model's explanatory power.

Consequently, we wish to solve the following program for the tremble q :

$$\min_q \left(M(q) - \Phi^{Na}(s, \mu, \mathbf{v}^{obs, s^0}) \right)^2. \quad (4.7)$$

Note, from the definition (4.6), that $M(\cdot)$ is a linear function of q . So we can solve program (4.7) by setting the derivative of the objective equal to zero. Compute from (4.6) that:

$$M'(q) = \sum_{\{(p,u)|0 < a^N(p,u) < 1\}} (1 - 2a^N(p,u)) \log \frac{a^N(p,u)}{1 - a^N(p,u)} \quad (4.8)$$

Now the f.o.c. for program (4.7) is:

$$2 \left(M(q) - \Phi^{Na}(s^0, \mu, \mathbf{v}^{obs, s^0}) \right) M'(q) = 0. \quad (4.9)$$

From (4.8), we see that generically, $M'(q) < 0$. (Each term in the sum in (4.8) is negative, except if $a^N = 1/2$.) Therefore, the solution of (4.9) requires:

$$M(q) = \Phi^{Na}(s^0, \mu, \mathbf{v}^{obs, s^0}). \quad (4.10)$$

Recalling that M is linear, we easily solve (4.10) for q , giving:

$$q_{\mu}^{*Nash} = \frac{\Phi^{Na}(s, \mu, \mathbf{v}^{obs, s^0}) - \left[\sum_{\{(p,u)|0 < a^N < 1\}} a^N \log a^N + \sum_{\{(p,u)|0 < a^N < 1\}} (1 - a^N) \log(1 - a^N) \right]}{\sum_{\{(p,u)|0 < a^N < 1\}} (1 - 2a^N) \log \frac{a^N}{1 - a^N}}. \quad (4.11)$$

Actually, this is the solution if the quantity on the r.h.s. of (4.11) lies in $[0,1]$. If the r.h.s. of (4.11) is greater than 1, then $q_{\mu}^{*N} = 1$ and if it is less than 0, then $q_{\mu}^{*N} = 0$. In other words, if the true solution of (4.9) were at a corner of $[0,1]$, the f.o.c. (4.10) becomes an inequality.¹¹

¹¹ It turns out that for all our countries and samples, the numbers $q^{*J}(s)$, for $J = \text{Nash, Kant}$ are in $(0,1)$. This means that, at the optimal values of q , $M(q^{*J}(s)) = \Phi^J(s, \mu, \mathbf{v}^{obs, s})$. We are able to adjust the tremble so that the expected value of the log likelihood of the trembling-hand model is precisely the *observed* log likelihood for that sample and model.

This completes the estimation procedure for the sample s^0 . We repeat the estimation procedure for each of $B=1200$ bootstrap samples. Denote, for bootstrap sample s , the q 's defined in equation (4.11) as $q^{*J}(s)$, for $J = \text{Nash, Kant}$.

We finally define two functions for all bootstrap samples s :

$$\Delta(s) = q^{*K}(s) - q^{*N}(s) \text{ and } \Gamma(s) = \Phi^K(s, \mu, \mathbf{v}^{obs,s}) - \Phi^N(s, \mu, \mathbf{v}^{obs,s}), \quad (4.12)$$

and deduce statistics on Δ and Γ using the 1200 bootstrap samples. For instance, if we find that the mean of the distribution $\Delta(s)$ is negative and more than two standard deviations below zero, we will say that the Kant model provides a better explanation of vaccination behavior than the Nash model, at the 95% significance level. A similar inference would be drawn if $\Gamma(s)$ is positive and at least two standard deviations away from zero.

In Section II of the Online Appendix, we perform a robustness check by running the program for several values of μ .

5 Major findings

We summarize the main findings of our analysis. The Online Appendix offers details on the survey, the bootstrap strategy, and the results for each country, as well as brief historical discussions of measles vaccination in each country.

- (i) For all countries, the profile of Kantian equilibrium strategies dominates the profile of Nash equilibrium strategies: that is, for all types (p, u) , $a^K(p, u) > a^N(p, u)$ or, equivalently, for all w , $\tilde{a}^K(w) > \tilde{a}^N(w)$. This is illustrated in two different spaces in Figure 1. Kantians always vaccinate with higher probability than Nashers. The continuous functions $\Delta\tilde{a}(w)$ are graphed in Figure 1a (recall the definitions in equations (2.17) – (2.19)). The observed values of w in any country sample comprise

a set of approximately 1000 values, which will lie along these curves. Figure 1b presents the graphs of the actual equilibrium profiles in the space $(a^N(p, u), a^K(p, u))$.

From Figure 1a, note that the differences between the Nash and Kantian strategies are greatest for Mexico: this is verified for the empirical distributions in the Mexican graph in Figure 1b. Contrast Mexico with Canada. We see from Figure 1a that the differences $\Delta\tilde{a}(w)$ are very small in Canada: this is verified in Figure 1b, where we see that observed strategy pairs are very close to (but lie above) the 45° line. We emphasize that the graphs in Figure 1a are derived from the estimated beta-distributions of types Q in the six countries.

- (ii) Our estimation procedure shows that the optimal tremble for the Kantian model, over all bootstrap samples, is significantly less than the optimal tremble for the Nash model. See Figure 2. In all six countries, the difference of the optimal trembles $(q^K - q^N)$ is significantly less than zero at the 99.9% significance level (that is, $\Delta(s) < 0$). Our interpretation of this fact is that the Kantian model provides a significantly better explanation of vaccination behavior than the Nash model, as we discussed in Section 4. In addition, in all six countries, the difference of the log likelihood functions $(\Gamma(s) = \Phi^K(s) - \Phi^N(s))$ is significantly greater than zero at the 99.9% significance level. See Figure 3. The interpretation is, again, that the Kantian model provides a significantly better explanation of vaccination behavior than the Nash model. The significance levels are reported in the Online Appendix.

In figures Figure 2 and Figure 3, we plot over each histogram the graph of the density function of the normal distribution with mean and standard deviation of the histogram, over the interval \pm three standard deviations from the mean, verifying visually our claim concerning significance levels.

- (iii) Figure 4 gives the histogram (over all bootstrap samples) of the values of the optimal tremble for the Nash and Kant model. These differ somewhat across countries.
- (iv) Figure 5 presents the histograms for the Nash and Kant log likelihood functions for all countries.
- (v) Figure 6 gives the histogram of the absolute value of the difference between the ‘observed’ coverage rate and the Nash (or Kant) equilibrium coverage rates at the optimal trembles, averaged over all samples. These differences are quite large. But recall that our estimation strategy is to minimize the tremble, not the difference between the equilibrium coverage rate and the observed coverage rate.
- (vi) Figure 7 and Figure 8 show that vaccine hesitancy with respect to measles vaccination is significantly less than hesitancy with respect to COVID-19 vaccination. We comment further on the distribution of vaccine hesitancy by political identification in the Online Appendix.




Fig. 1a. The difference between the Kantian and the Nash equilibrium strategy profiles: $\Delta \tilde{a}(w)$, where $w = p(1 - u)$.




Fig. 1b. Equilibrium strategy profiles from the original survey data. Kantian vs. Nash strategies for the observed strategy pairs $(\tilde{a}^N(w), \tilde{a}^K(w))$

Figure 1. Kantian vs. Nash strategy profiles.

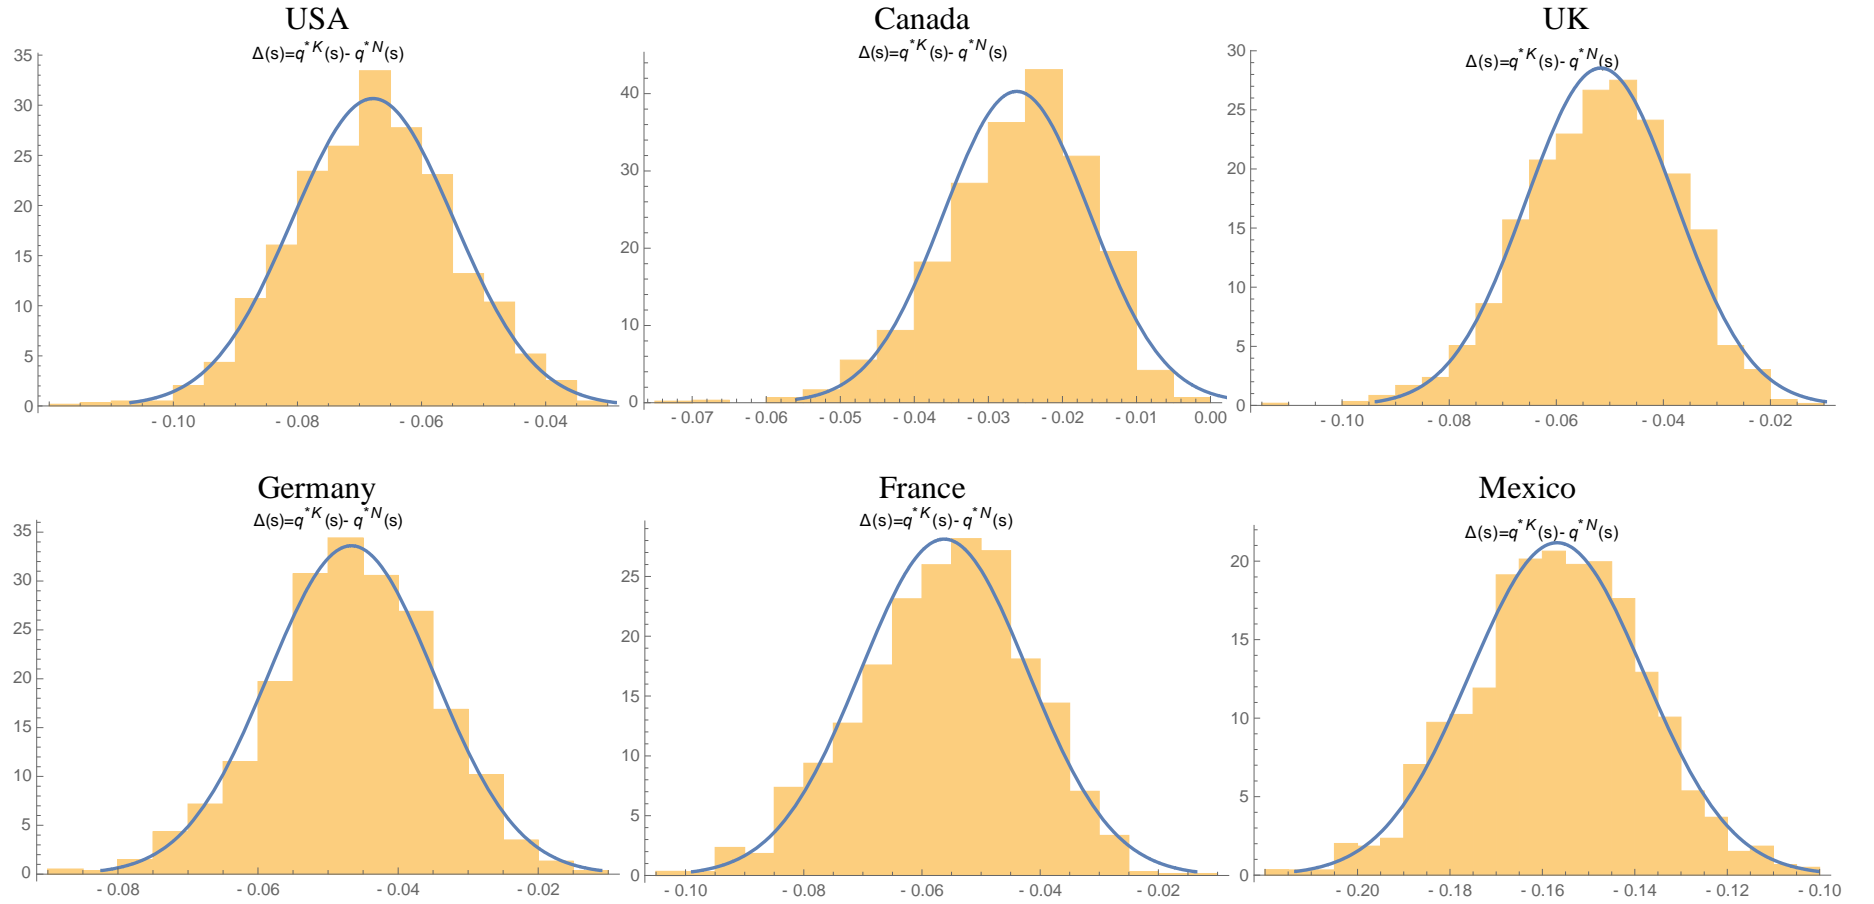


Figure 2. Probability density histogram of $\Delta = q^{*K} - q^{*N}$, and the PDF of a Normal distribution $N(m, \sigma)$ with $m = \text{Mean}(\Delta)$ and $\sigma = \text{StaDev}(\Delta)$, truncated at three standard deviations from the mean.

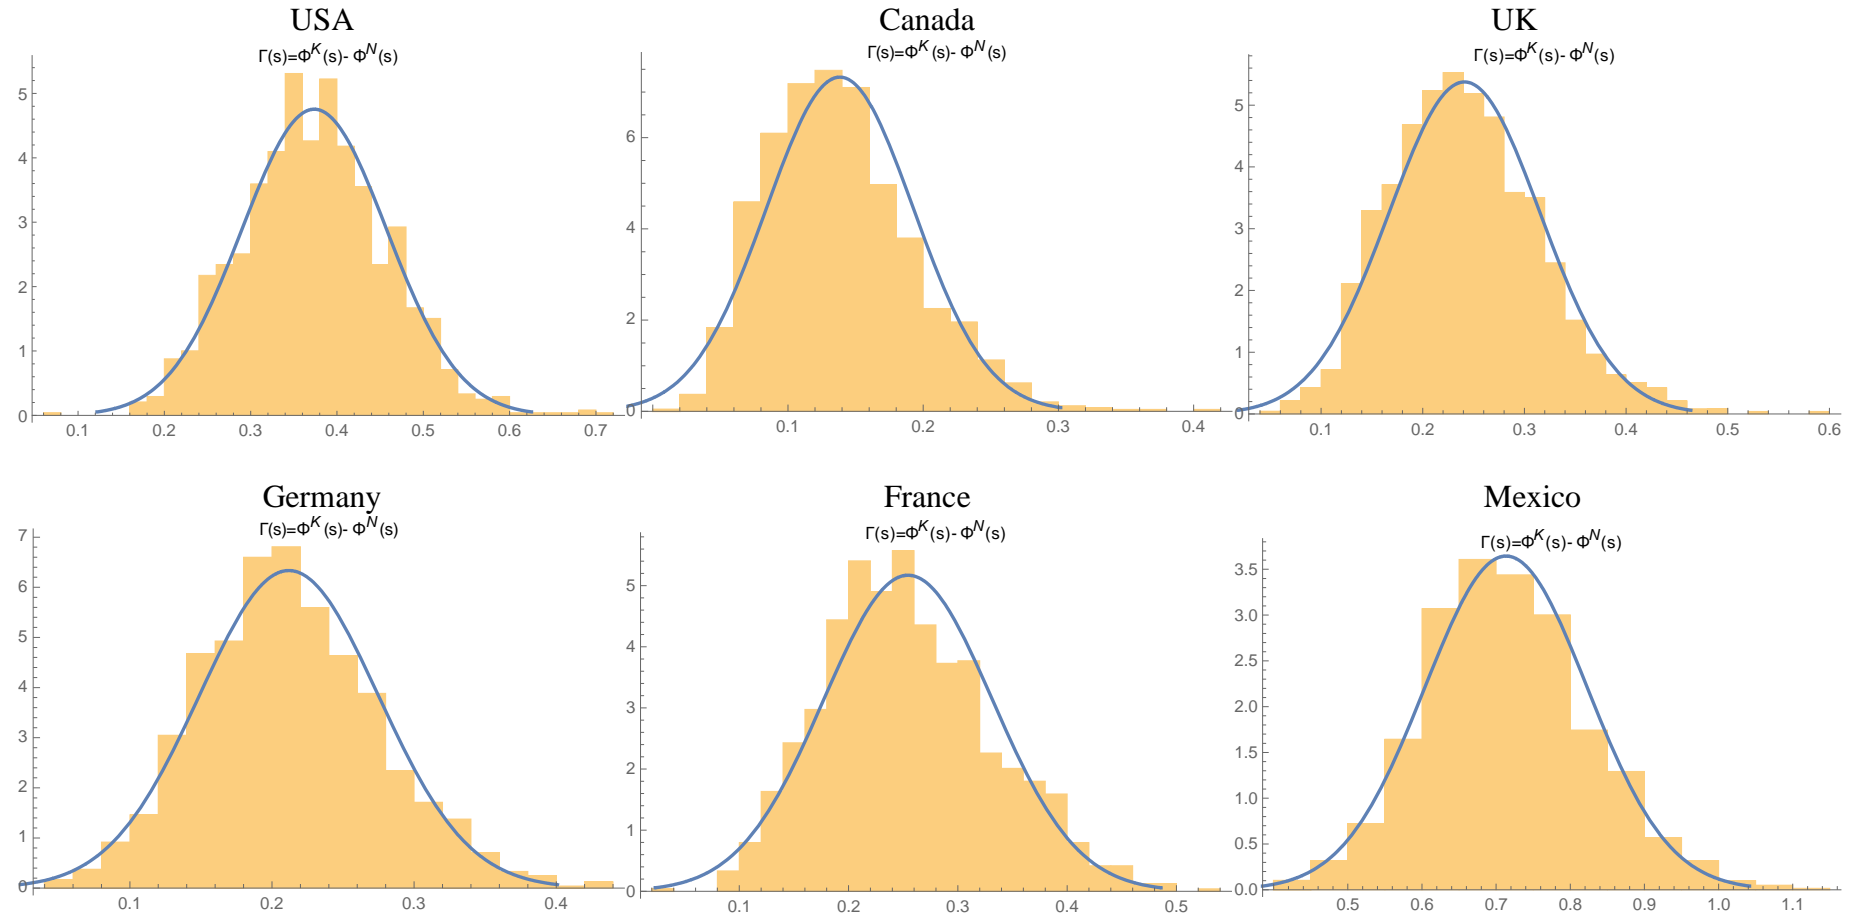


Figure 3. Probability density histogram of the difference of the per capita log-likelihoods ($\Gamma = \Phi^N - \Phi^K$), and the PDF of a normal distribution $N(m, \sigma)$ with $m = \text{Mean}(\Gamma)$ and $\sigma = \text{StaDev}(\Gamma)$, truncated at three standard deviations from the mean.

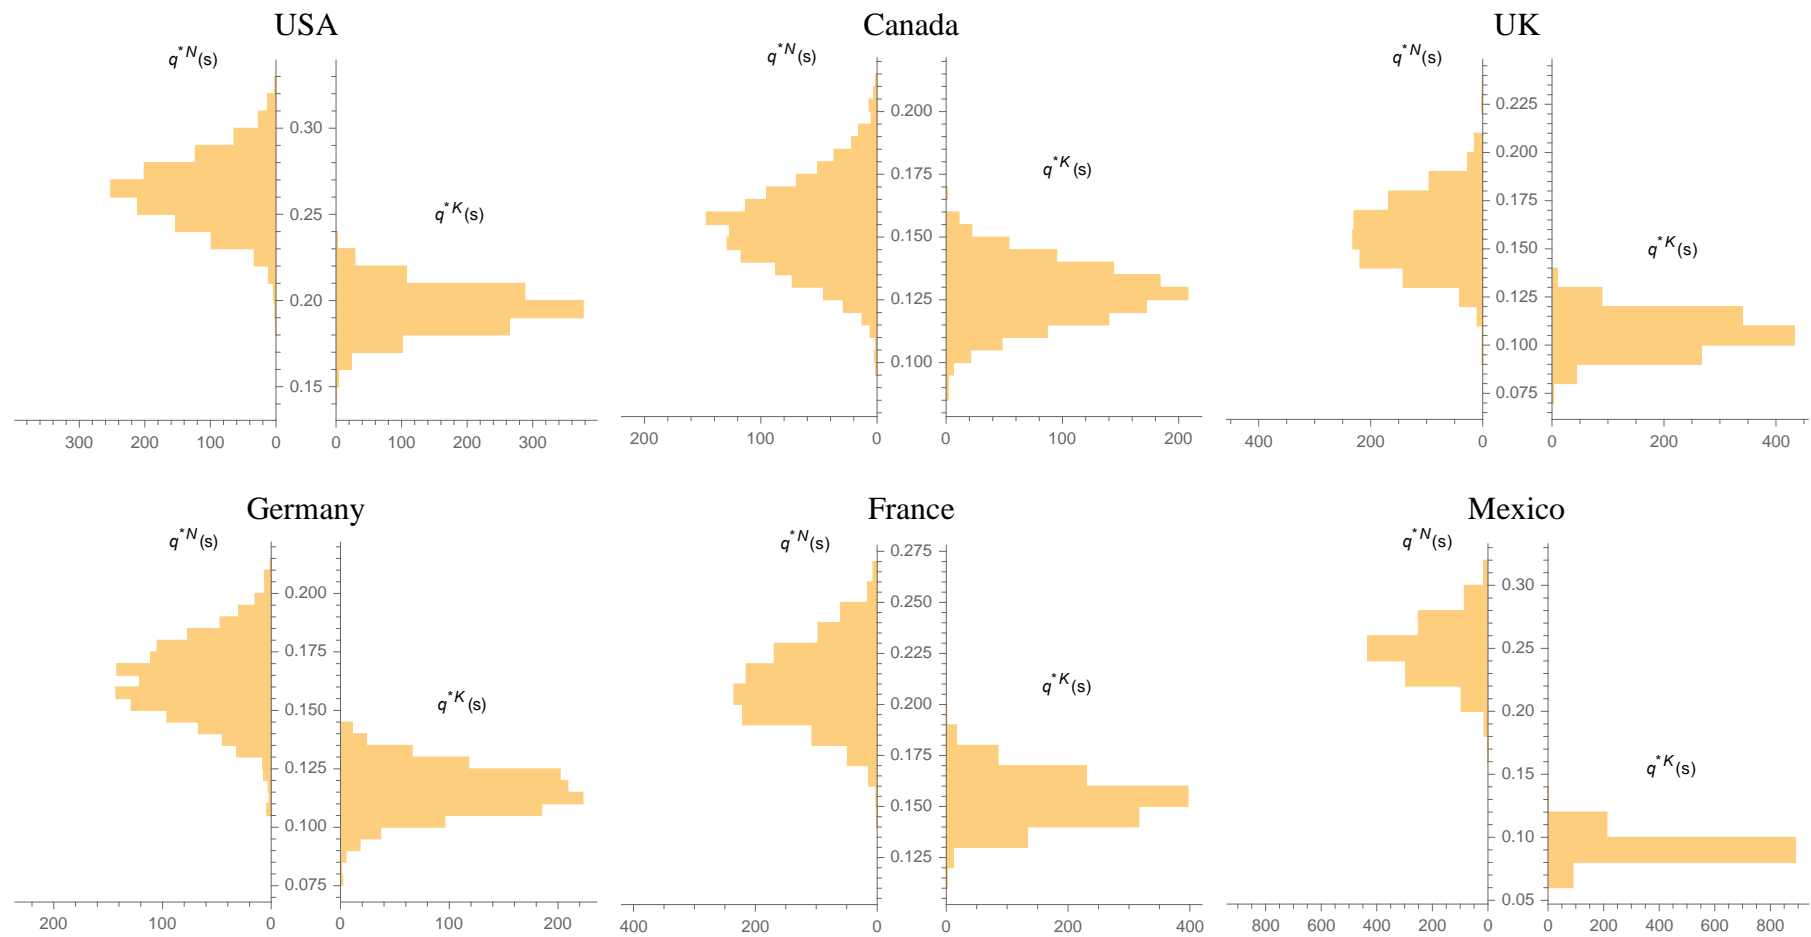


Figure 4. Histograms of the optimal tremble for the Nash (q^{*N}) and Kantian (q^{*K}) equilibrium.

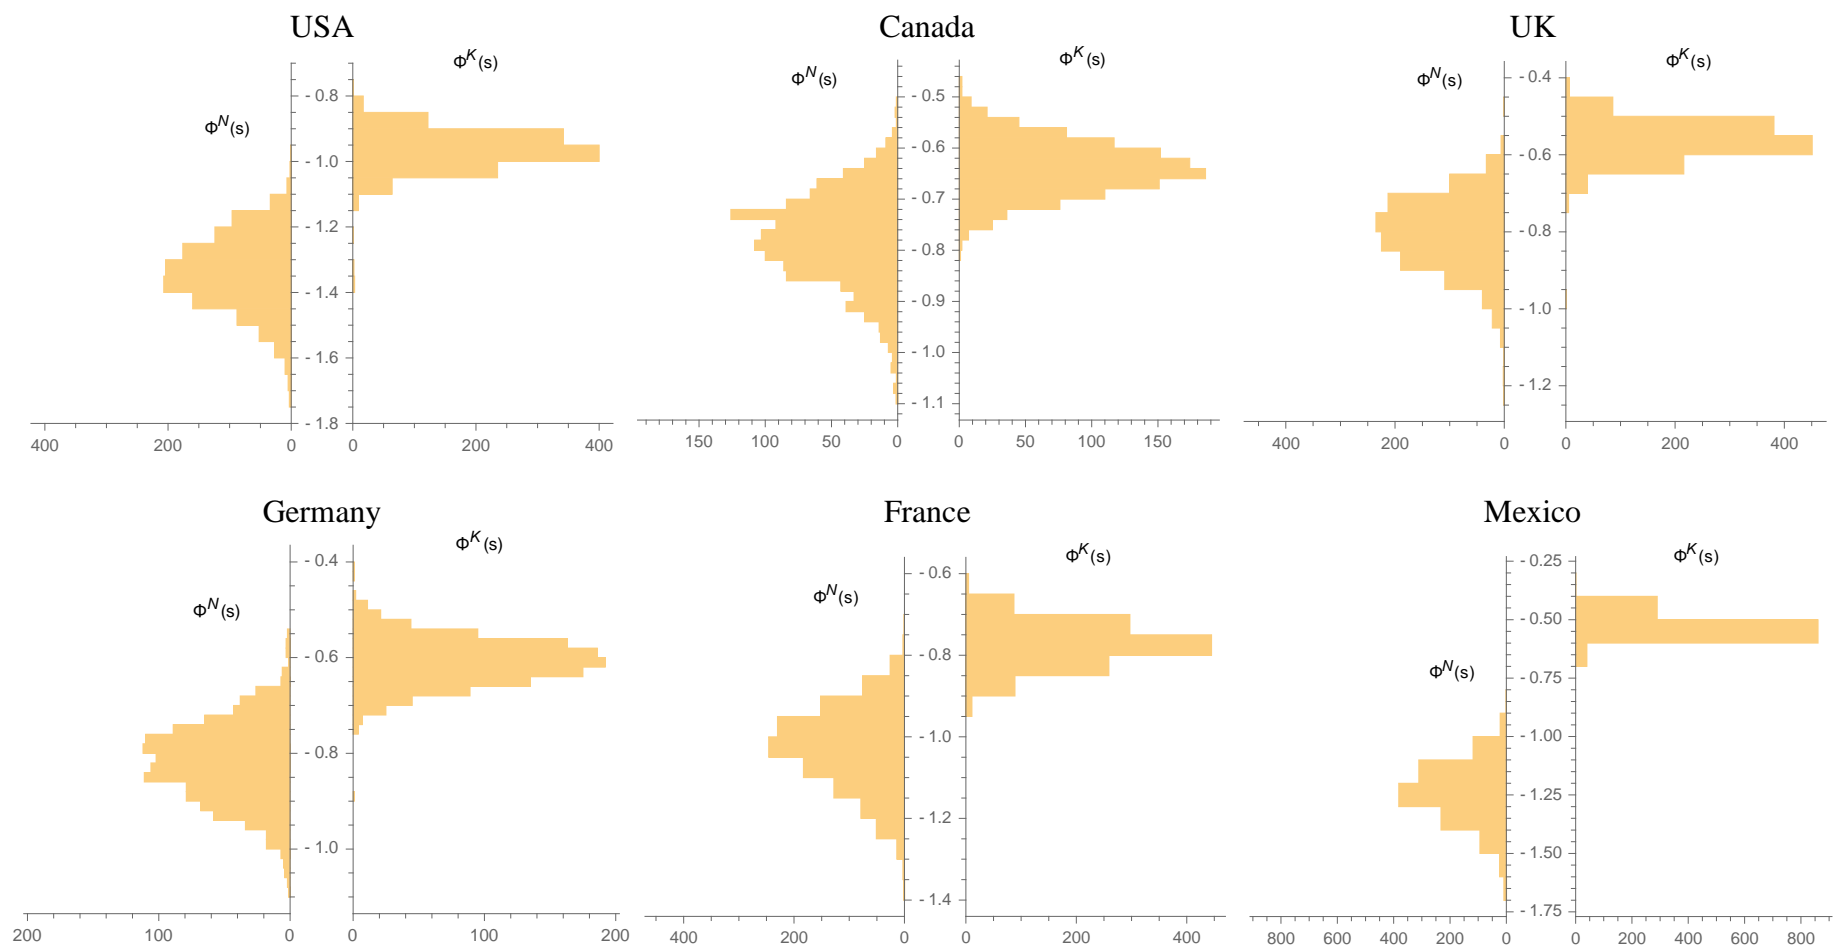


Figure 5. Histograms of the Nash (Φ^N) and Kantian (Φ^K) log-likelihoods.




Figure 6. Distributions of the difference between coverage at the survey sample (\bar{a}_0) and at the Nash and Kantian equilibrium coverage ($\bar{a}^J, J = N, K$) at the 'optimal' tremble.




Figure 7. Percentage of respondents who report choosing to vaccinate against measles among different ideological groups. Vaccinating against measles means that respondents either vaccinated their child or would have vaccinated their child had they had one. Ideological identification is self-reported.¹⁸ “Other” includes those who responded “not sure” and no answers. Surveys conducted in December 2020.




Figure 8. Percentage of respondents who would vaccinate against Covid-19 if a vaccine were available at no cost. Ideological identification is self-reported.¹⁹ “Other” includes those who responded “not sure” and no answers. Surveys conducted in December 2020.

¹⁸ Ideological labels respond to the options offered in the survey for all countries but France, where respondents had to locate in a Left-Right dimension with a 0 to 10 scale. For comparison purposes, responses were recoded as follows: Very Liberal = {0,1}, Liberal = {2,3}, Moderate = {4,5,6}, Conservative = {7,8}, and Very Conservative = {9,10}.

¹⁹ See previous footnote.

6 Why does the Kantian model of vaccination give a superior explanation?

We propose that the reason the Kant model is superior to the Nash model is that it gives uniformly higher probabilities of vaccination than the latter, and the coverage rates in the Kantian equilibria are closer to the observed coverage rates in our samples than the Nash coverage rates. See the six panels in Figure 9. Coverage rates at the survey sample (\bar{a}_0), at the Nash equilibrium strategies ($\bar{a}^N(q^{*N})$), and at the Kantian equilibrium strategies ($\bar{a}^K(q^{*K})$). For all countries, $\bar{a}^N(q^{*N}) < \bar{a}^K(q^{*K})$. For all countries, except for Canada and the UK, $\bar{a}^N(q^{*N}) < \bar{a}^K(q^{*K}) < \bar{a}_0$. See the Online Appendix for numerical comparisons, where we plot the equilibrium coverage rates in the Nash and Kant equilibria of our 1200 bootstrap samples, in comparison to the observed coverage rate of sample s^0 (for each country). We believe that actual vaccination behavior is more pervasive than the Nash model predicts because in all countries in our sample there is a social norm to vaccinate. The social norm shares with the Kantian equilibrium the property of inducing more pervasive vaccination behavior than Nash equilibrium predicts. Even though the Kantian equilibrium strategies appear to be not much larger than the Nash probabilities, the fact that they are *always* larger than the latter, on the space of types, makes the likelihood of the Kantian equilibrium significantly greater than the likelihood of the Nash equilibrium. (See our major finding (i) in Section 5.)

Why has a social norm to vaccinate developed in all these countries? Because, we conjecture, of the phenomenon of herd immunity. We pointed out in Section 1 that the Nash equilibrium of the vaccination game will always be Pareto inefficient, because of the positive externality associated with vaccination. Kantian optimization is *one way* to repair this inefficiency: *another way* to do so is through the development of a social norm to vaccinate. We are suggesting that the latter has indeed occurred in many societies, and this is why the Kantian equilibrium gives a better explanation of vaccination behavior than the Nash equilibrium.

We might point out that a similar explanation applies to explaining the payment of income taxes, which, in all highly developed countries is more pervasive than Nash equilibrium predicts. (There is a literature

showing that in the United States, given the penalties associated with tax evasion and the probability of being caught for doing it, evasion should be significantly higher than it is, assuming reasonable attitudes towards risk (see Alm (2019) and the many references therein). The reason is that there is a social norm to pay one's taxes. And this norm has developed precisely because the 'tax game' is a monotone increasing game: the taxes you and others pay make me better off, because they finance public goods and social insurance from which I benefit.²⁰

To test this conjecture, we administered a second survey on the motivation to vaccinate, which was inadequately covered in our first survey, in the US and France, two countries that have significant anti-vax movements. We report the key findings of both surveys in the next three tables. In our follow-up survey we received 1243 responses from Americans and 1490 responses from French residents.²¹ Herd immunity and vaccination behavior of others is clearly indicated to be encouraging rather than discouraging own vaccination, which is in line with Kantian optimization and a strong social norm, but not with Nash equilibrium. (A Nash optimizer will be discouraged to vaccinate her child if herd immunity is attained.) In a scenario of well-established herd immunity for a child illness, 68.4% of US respondents (57.5% of French ones) are either 'strongly encouraged or encouraged' to vaccinate their own child (Table 5), while only 6.24% (7.04%, resp.) are discouraged.

²⁰ Or perhaps the tax norm has developed for another reason: that people have an aversion to cheating.

²¹ These counts include all responses (including those who simply declined consent and ended the survey) but exclude any responses classified as "spam" by Qualtrics. We report in Appendix B the results obtained without removing those respondents who either did not answer all questions (for both the US and France) or gave low quality responses (US only). Results are qualitatively similar.

Table 5. “Imagine herd immunity is already well-established for a specific child illness because of a high vaccination rate. Would that encourage or discourage you from vaccinating your own child? (Q4.6; US and France)”

	US		France	
	Frequency	Percent	Frequency	Percent
Strongly encourage	475	43.58	372	28.48
Encourage	271	24.86	379	29.02
Leave unchanged	276	25.32	463	35.45
Discourage	41	3.76	58	4.44
Strongly discourage	27	2.48	34	2.60
N	1,090	100	1,306	100

The same reaction is observed to an individual act of vaccination. Learning that others have vaccinated their child ‘strongly encourages or encourages’ 61.4% of US respondents (and 45.9% of French ones, see Table 6). Likewise, own vaccination is expected to strongly encourage or encourage others’ vaccination by 64.6% of US respondents (and 50.4% of French ones, see Table 7), indicating that a parent believes a social norm is operative.

Table 6. “If you learn that others have vaccinated their child, would that encourage or discourage you to vaccinate your child? (Q4.3; US and France)”

	USA		France	
	Frequency	Percent	Frequency	Percent
Strongly encourage	390	35.78	236	18.02
Encourage	279	25.60	365	27.86
Leave unchanged	368	33.76	673	51.37
Discourage	25	2.29	15	1.15
Strongly discourage	28	2.57	21	1.60
N	1,090	100	1,310	100

Table 7. “When you vaccinate your child, would you expect others to be encouraged or discouraged by your action to also vaccinate their child? (Q4.4; US and France)”

	USA		France	
	Frequency	Percent	Frequency	Percent
Strongly encouraged	360	33.03	255	19.47
Encouraged	344	31.56	405	30.92
Leave unchanged	351	32.20	624	47.63
Discouraged	19	1.74	16	1.22
Strongly discouraged	16	1.47	10	0.76
N	1090	100	1310	100

For those respondents who have or would vaccinate their child, 73.5% of US respondents (and 72.6% of French ones) indicate that “Vaccination protects my child from disease” is a very important reason for that decision. Other reasons that relate to herd immunity or social norms are also deemed very important by large fractions of respondents, “Vaccination of my child contributes to herd immunity” by 54.9% of US respondents (46.2% of French ones) and “I vaccinate because other parents I know choose to vaccinate” by 35.4% of US respondents (20.4% of French ones). Conversely, for those respondents who have not or would not vaccinate their child, side effects and choice autonomy are deemed as very important more often than matters of herd immunity or social norms. This can be seen in the contrast between “There are possibly severe side effects to vaccination” (56.4% in the US, 61.6% in France) and “Vaccination should be a matter of free choice” (54.3% in the US, 58.0% in France) on the one side and on the other side “If vaccination coverage is already high in the community, my child will be safe without vaccination” (30.9% in the US, 24.1% in France) and “Other parents I know are choosing not to vaccinate” (31.9% in the US, 32.1% in France).




Figure 9. Coverage rates at the survey sample (\bar{a}_0), at the Nash equilibrium strategies ($\bar{a}^N(q^{*N})$), and at the Kantian equilibrium strategies ($\bar{a}^K(q^{*K})$). For all countries, $\bar{a}^N(q^{*N}) < \bar{a}^K(q^{*K})$. For all countries, except for Canada and the UK, $\bar{a}^N(q^{*N}) < \bar{a}^K(q^{*K}) < \bar{a}_0$. See the Online Appendix for numerical comparisons.

7 Conclusion

The vaccination of children can be modelled as a game in which there are significant positive externalities from the individual's choice to vaccinate; we say the game is monotone increasing. The Nash equilibria of such games are inefficient, a fact colloquially known as the free-rider problem. The Kantian equilibria of such games are generically efficient. We have shown, in a sample of six countries, that the Kantian model provides a superior explanation of vaccination behavior than the Nash model.

This might seem surprising: can we argue that some social institution is guiding parents to optimize in the Kantian manner? We do not do so here. However, we do conjecture that a fairly strong social norm is operative in many countries that induces parents to vaccinate at higher rates than they would were they playing the Nash equilibrium of the game. The social norm shares this property with the Kantian equilibrium, which, we argue, is the reason the Kantian model is more predictive of real behavior than the Nash model.

Indeed, the same reasoning applies to monotone *decreasing* games –ones in which each player's contribution imposes a negative externality on the welfare of others. The Nash equilibrium in monotone decreasing games is also inefficient, a fact known colloquially as the tragedy of the commons. The canonical fishing game, in which an increase in each fisher's labor decreases the productivity of the fishery for other fishers, is the textbook example. Elinor Ostrom (1990) studied common-pool resource problems in many communities around the world, and found that the participants often achieve efficient solutions, in the face of Nash equilibria that suffer from the tragedy of the commons, through the evolution of social norms that induce them to internalize the externalities of their behavior. Unfortunately, when Ostrom worked, the Kantian tool was not yet available: we suggest its fruitful application to the problems she studied.

APPENDIX A: Estimation of the parameter γ

We use the source “WHO vaccine-preventable diseases: Monitoring system, 2020 global summary,” https://apps.who.int/immunization_monitoring/globalsummary/, which contains data for a large set of countries on infectious disease immunization rates and morbidity.

A *cohort* of children is the set of children in the country born in a given year.

For a particular country, let:

n' = total population of children ages 0-5 in year t , $t = 2015, \dots, 2019$

r' = measles immunization coverage rate, children under 5, year t

c' = number of measles cases, year t

\bar{u} = number of susceptible children under 5 in a given cohort

$\bar{n} = \frac{\sum_{t=1}^5 n^t}{5}$; $\bar{n}/5$ = number of children in a given cohort

$\bar{r} = \frac{\sum_{t=1}^5 r^t}{5}$

$\bar{c} = \frac{\sum_{t=1}^5 c^t}{5}$

p = probability that a susceptible child of a given cohort contracts measles in a given year

π = probability that a susceptible child of a given cohort contracts measles by five years of age

By definition, $\bar{u} = \frac{\bar{n}}{5} (1 - \bar{r})$. The median age of contracting measles is age five. Therefore, the

number of cases of measles of children under five in a given cohort in a given year is $\frac{\bar{c}}{10}$. Therefore $p =$

$$\frac{\bar{c}/10}{\bar{u}} = \frac{\bar{c}}{10\bar{u}}.$$

Assume that an unvaccinated (susceptible) child in a given cohort has a probability p of contracting measles in each year under five. Then:

$$\pi = p + p(1 - p) + \dots + p(1 - p)^4 = p \frac{1 - (1 - p)^5}{p} = 1 - (1 - p)^5.$$

In our model we have $\pi(r) = (1 - r)^\gamma$. We propose that $\pi(r)$ is precisely the value π defined above: as a parent, I am concerned with the probability that my young child contracts measles if I choose not to vaccinate her, knowing that the coverage rate is r .

As described in the text, we assume the contagion process in North America (Canada and the US) is different from in Europe (UK, Germany, France). For each country j , we compute a data point (r, π) . Hence, we compute two values of γ : γ^{NA} gives the best fit of the function $\pi(\cdot)$ to the points $\{(r^{\text{US}}, \pi^{\text{US}}), (r^{\text{Can}}, \pi^{\text{Can}})\}$ and γ^{EUR} gives the best fit of the function $\pi(\cdot)$ to the points $\{(r^j, \pi^j) | j \in \{\text{UK, France, Germany}\}\}$. See figures Figure A.1 and Figure A.2.

Unfortunately, the data set does not provide measles morbidity for Mexico.




Figure A.1 Fitting the function $\pi(\cdot)$ for the US and Canada: $\gamma^{\text{NA}} = 3.110$.




Figure A.2 Fitting the function $\pi(\cdot)$ for the four European countries (Sweden included): $\gamma^{\text{EUR}} = 1.995$.

Appendix B: Survey Two with all responses

Here we reproduce the tables from Section 6 without removing those respondents who either did not answer all questions (for both the US and France) or gave low quality responses (US only). The total N for each question includes only those respondents who answered one of the five options — “blank” (no response) answers are excluded from the total N and percentage calculations.

Table B.1 “Imagine herd immunity is already well-established for a specific child illness because of a high vaccination rate. Would that encourage or discourage you from vaccinating your own child? (Q4.6; US and France)”

	US		France	
	Frequency	Percent	Frequency	Percent
Strongly encourage	475	43.58	372	28.48
Encourage	271	24.86	379	29.02
Leave unchanged	276	25.32	463	35.45
Discourage	41	3.76	58	4.44
Strongly discourage	27	2.48	34	2.60
N	1090	100	1306	99.99

Table B.2 “If you learn that others have vaccinated their child, would that encourage or discourage you to vaccinate your child? (Q4.3; US and France)”

	US		France	
	Frequency	Percent	Frequency	Percent
Strongly encourage	390	35.78	236	18.02
Encourage	279	25.60	365	27.86
Leave unchanged	368	33.76	673	51.37
Discourage	25	2.29	15	1.15
Strongly discourage	28	2.57	21	1.60
N	1090	100	1310	100

Table B.3 “When you vaccinate your child, would you expect others to be encouraged or discouraged by your action to also vaccinate their child? (Q4.4; US and France)”

	US		France	
	Frequency	Percent	Frequency	Percent
Strongly encouraged	360	33.03	255	19.47
Encouraged	344	31.56	405	30.92
Leave unchanged	351	32.20	624	47.63
Discouraged	19	1.74	16	1.22
Strongly discouraged	16	1.47	10	0.76
N	1090	100	1310	100

For the other questions the results are also very similar to those in the main text when including all responses.

Table B.4 Percentage of respondents who have or would vaccinate their child(ren) answering that a particular reason is “very important” in their choice to do so.

	US	France
Vaccination protects my child from disease	73.49%	72.57%
Vaccination of my child contributes to herd immunity	54.84%	46.24%
I vaccinate because other parents I know choose to vaccinate	35.38%	20.42%

Table B.5 Percentage of respondents who have not or would not vaccinate their child(ren) answering that a particular reason is “very important” in their choice to not do so.

	US	France
There are possibly severe side effects to vaccination	56.38%	61.61%
Vaccination should be a matter of free choice	54.26%	58.04%
If vaccination coverage is already high in the community, my child will be safe without vaccination	30.85%	24.11%
Other parents I know are choosing not to vaccinate	31.91%	32.14%

References

- Alm, J. 2019. “What motivates tax compliance?.” *Journal of Economic Surveys* 33, 353-388.
- Avery, C., W. Bossert, A. Clark, G. Ellison and S.F. Ellison, 2020. “An economist’s guide to epidemiology models of infectious disease,” *Journal of Economic Perspectives*. 34, 79-104.
- Dabbagh, A., R.L. Laws et alii. 2018. “Progress toward regional measles elimination –Worldwide, 2000-2017,” *Morbidity and Mortality Weekly Report* 67 (47), US Dept. of Health and Human Services, Center for Disease Control and Prevention.
- Holt, C.A. and S.K. Laury, 2002. “Risk aversion and incentive effects,” *American Economic Review* 92, 1644-1655.
- Ostrom, E. 1990. *Governing the commons: The evolution of institutions for collective action*, Cambridge, UK: Cambridge University Press.
- Roemer, J.E. 2019. *How we cooperate: A theory of Kantian optimization*, New Haven: Yale University Press.
- Strebel, P. M. and W. A. Orenstein. 2019. “Measles,” *New England Journal of Medicine*. 381(4), 349–357. PMID: 31184814.