

Integrated epi-econ assessment of vaccination^{*}

Timo Boppart[†], Karl Harmenberg[‡],
Per Krusell[§] and Jonna Olsson[¶]

January 11, 2022

Abstract

Using an integrated epi-econ model, we compute the value of vaccines for Covid-19, both under a planner's solution and in competitive equilibrium. The specific model, developed in Boppart, Harmenberg, Hassler, Krusell, and Olsson (2021), factors in not just value-of-life aspects along with standard economic variables but also the value of leisure activities that rely on a social component. We find that the societal value of vaccination is large; we estimate that, translated into monetary terms, the value of vaccinating one young individual in the competitive equilibrium is \$17,800. Externalities are large: less than half the societal value is internalized by individuals (assuming that they act purely in their self-interest). Finally, behavioral responses are important, with a substantial share of the value of vaccines being attributed to people enjoying more socially-oriented leisure when more people are vaccinated.

^{*}We thank our discussant, Marina Azzimonti, and conference participants at Fed St. Louis-JEDC-SCG-SNB-UniBern Conference for helpful comments.

[†]Institute for International Economic Studies, University of St. Gallen

[‡]BI Norwegian Business School

[§]Institute for International Economic Studies

[¶]University of Edinburgh

1 Introduction

It behooves economists to produce frameworks and analyses that can speak to policy questions that involve tradeoffs between different economic variables. When life and death are added to the consequences of policy choices, one might think that we are moving outside the realm of economics. Our view, however, is rather the opposite: then it appears even more important to make the tradeoffs clear and produce coherent comparisons using explicit models. One such instance is Covid-19. In earlier work, Boppart, Harmenberg, Hassler, Krusell, and Olsson (2021), we constructed an *epi-econ assessment model* with this aim and used it to examine the performance of an unregulated market economy (pitched against a well-defined social planning outcome) under various kinds of epidemics. In the present paper, we use this framework to evaluate the efficacy, and value, of vaccines, i.e., we consider vaccination an explicit choice. This allows us not only to assess vaccines from a positive perspective and compute their monetary equivalent values but also to see what the relative values are of vaccinating different subgroups of the population. In our simple model, we distinguish “young” from “old” agents; aside from differing in typical ways (life expectancy, productivity), these two groups also differ in their Covid-19 vulnerability.

The integrated epi-econ assessment framework we need for policy evaluation, in our view, needs to not only include explicit epidemiological modeling but also explicit economic decision such that one can distinguish market outcomes—where presumably at least parts of the population behave only in a self-interested way—from a well-defined planning problem. The economic structure, moreover, should lay out how people’s time allocation affects both economics (production and leisure) and epidemiology (epidemic spread through social interaction). It should model how epidemic risk affects economic decision making and how economic actions affect the effective reproduction number. When welfare is concerned, it should factor in aspects of people’s lives such as life expectancy and also include ways in which social activities—contacts with others—matter for people.

We also need a framework that is quantitative: qualitative insights are simply not useful enough in this important policy area. Of course, not all features of reality can be included in an integrated assessment model and, surely, we abstract from some issues here that we believe are important to take into account; we discuss these in our concluding remarks (in Section 5). But for the features that are present in the model, we discipline the selection of functional forms and parameters with the use of data. For example, for one of the features that distinguish our work from that of others—the incorporation of explicit time allocation problems that allow consumers to choose more

or less social interaction—we use data from ATUS (American Time Use Survey).

Before describing the results, let us very briefly mention a few important features of our setting. We assume a population structure where each individual belongs to a (large) family within which there is full insurance against any shocks; at the same time, there are many such families. This assumption is admittedly somewhat extreme, but we know from the vast literature on macroeconomic and inequality that significant insurance is obtained by precautionary saving so we regard our framework at least as a useful starting point. A benefit of this assumption, i.e., the presence of a “representative family”, is that our social planner’s welfare criterion—which is also used to evaluate vaccines—naturally becomes that of the representative family. Second, the way we introduce different extents of social interactions is by indexing goods, and leisure activities involving goods consumption, by their respective degrees of social interaction. We dichotomize: goods (and services) are either consumed in public or in private, where in-private consumption does not involve social interaction. Similarly, workers can contribute to production either from home or by working “in the office”. We assume utility functions over different goods and leisure activities that involve non-trivial substitution and, similarly, there is nontrivial substitution between work at home and working in the office, all constrained to match available microeconomic data. Third, the epidemiological model is a simple SIR framework with a mortality risk that involves some crowding at hospitals. All these assumptions are identical to those employed in our earlier, main paper.

The key take-aways from our analysis are as follows. The representative family would pay \$7,800 to vaccinate a young individual in the beginning of the epidemic, and \$9,600 to vaccinate an old (who is more vulnerable to the disease). That is, from the perspective of a self-interested family, investing money in vaccinating its old members is worth more than vaccinating its young members. However, for either kind of vaccination, the value is very large. The societal value of a vaccine dose, i.e., that which takes into account the effect on the lower spread of the virus (we assume that the vaccinated, like the “recovered”, do not spread the virus), (a) is much higher still, and (b) the order of prioritization across age groups is the reverse: the benefit of vaccinating a young agent is \$17,800 whereas that of vaccinating an old agent is \$11,600). A large component of this difference is that the young are more socially active and thus more intense spreaders of the virus (which constitutes the main externality in our framework). A corollary of this fact is that it is very important to take into account how different agents systematically interact differently. These numbers apply to a marginal vaccination: to the vaccination of a single household member. If a

significant mass of the population is vaccinated, of course the marginal effects change. If a large chunk of the *young* population is vaccinated, there are substantial benefits for the *old* population; 40% of total benefits in such a scenario arise from the old being able to enjoy more social leisure during the course of the epidemic.

In terms of related literature, there are of course many papers by now that examine, from various perspectives, the interplay of macroeconomics and epidemiology. We survey this literature in our earlier paper Boppart, Harmenberg, Hassler, Krusell, and Olsson (2021). The literature on the economics of vaccination is considerably smaller. An early theoretical contribution is Chen and Toxvaerd (2014). Arellano et al. (2021) study the effect of delayed vaccine distributions on emerging markets with financial frictions and default risk. Auld and Toxvaerd (2021) document that the roll-out of Covid-19 vaccines coincided with fewer infections and more social activity, which is consistent with our model. Goodkin-Gold et al. (2020) show, in an analytically tractable setting, that the vaccine externality is non-monotonic in the reproduction rate of the disease. Overall, our setting is considerably richer than these papers. In this special issue of the journal, Glover et al. (2021) and Garriga et al. (2021) also study the value of vaccination in integrated economic-epidemiological model. Here, Glover et al. (2021), whose setting is relatively close to ours, also include age heterogeneity in their model; interestingly, they find that the societal value of vaccinating the old exceeds that of vaccinating the young. Despite their similarities, the models are sufficiently different that it is difficult to pin down the exact sources of this difference. We suspect, however, that our result is a function of our different modeling of social interactions. Finally, and interestingly, there are attempts to develop theory of social interactions along the dimension of social networks, which appear important in this context; see Azzimonti et al. (2020) for a prominent example. Our framework does not have social interaction for the sake of social interaction—social interaction always occurs while consuming or producing specific goods and services—but it seems possible to extend our framework in that direction.

The paper is organized as follows. In Section 2 we very briefly describe our model; we do not spend much time on it, however, mostly referring the reader to the original paper. In Section 3, we discuss the calibration, and in Section 4 we look at the results. Section 5 concludes, in particular placing emphasis on those among the features missing in our model that we believe are most important.

2 An epi-econ framework

We study vaccination using the model from Boppart, Harmenberg, Hassler, Krusell, and Olsson (2021). In this section, we provide a brief summary of the model framework.

There are two types of individuals: young and old. Both types have a choice of how to spend their time: they can work, or they can enjoy leisure. The leisure can be spent in two different sectors: one that we label as the in-priVate (V) sector, and one that we label as the in-puBlic (B) sector. When people spend their time on in-priVate leisure they are on their own (e.g., they are at home watching television). In other words: they do not interact with others and hence there is no risk of getting infected. The goods and services used for the in-priVate leisure are produced in the workplace or at home. When people work in the workplace, they interact with their colleagues, and there is a risk of spreading the virus.

In the in-puBlic sector, consumption and work take place jointly, and the virus can be spread between those enjoying their leisure (e.g., customers in the restaurant) and those working in the sector (e.g., waiters in the restaurant). However, even in this sector there is a possibility (for at least some employees, e.g., the restaurant's accountant) to work from home without physically interacting with others.

In the beginning of the epidemic, a small fraction of the population is infected, while the majority are still susceptible, and no one has yet recovered. As the epidemic spreads, susceptible individuals become infected, while some infected individuals recover and a fraction of the infected dies. The infection fatality rate is higher for the old than for the young, and it increases if the hospitals become overcrowded. We assume full immunity, so that once an individual has recovered, he/she cannot be re-infected. Individuals in the model receive utility from leisure, from consumption, and from the intrinsic value of being alive.

2.1 Recursive formulation of the planner's problem

We denote the complete SIR state by Epi, so that $\text{Epi} = [S^y, I^y, R^y, S^o, I^o, R^o]$. S_t^i is the number of susceptible of type $i \in \{y, o\}$ (young/non-vulnerable and old/vulnerable, respectively) at time t ; similarly I_t^i is the number of infected and R_t^i the number of recovered. The total population size at time t is $\sum_i (S_t^i + I_t^i + R_t^i)$, with initial conditions satisfying $I_0^y + I_0^o = 0.001$, $R_0^i = 0$ for $i \in \{y, o\}$ and $S_0^y + S_0^o = 0.999$. We assume that the fraction initially infected is equal in the young and the old population. In the absence of an epidemic, all young individuals live until period $\mathcal{T}^y$ and all old individuals live until period $\mathcal{T}^o < \mathcal{T}^y$.

The epidemic is over at time $\mathcal{T}$ which is chosen to be large enough so that the epidemic is endogenously finalized. The planner chooses how much time the young and old should spend working in the workplace in respective sector (n_{Bw}^i and n_{Vw}^i), how much time to spend working at home (n_{Bh}^i and n_{Vh}^i), how much time to spend on leisure in the respective sector (h_B^i and h_V^i), and how to allocate consumption in the two sectors (c_B^i and c_V^i) up until time $\mathcal{T} - 1$ and maximizes:

$$V_t(\text{Epi}) = \max \left\{ \sum_i (S^i + I^i + R^i) v(c_B^i, h_B^i, c_V^i, h_V^i) + \beta V_{t+1}(\text{Epi}') \right\} \quad t < \mathcal{T}$$

subject to non-negativity constraints for time and consumption quantities, the resource constraints (using the short-hand notation $\phi^i = S^i + I^i + R^i$),

$$\begin{aligned} F_B(\phi^y n_{Bh}^y, \phi^y n_{Bw}^y, \phi^o n_{Bh}^o, \phi^o n_{Bw}^o) &= \sum_i \phi^i c_B^i, \\ F_V(\phi^y n_{Vh}^y, \phi^y n_{Vw}^y, \phi^o n_{Vh}^o, \phi^o n_{Vw}^o) &= \sum_i \phi^i c_V^i, \end{aligned}$$

the time-allocation constraints (for both young and old),

$$1 = h_B^i + h_V^i + n_{Bh}^i + n_{Vh}^i + n_{Bw}^i + n_{Vw}^i,$$

the evolution of the epidemic (for both young and old),

$$\begin{aligned} T^i &= (\hat{\pi}^B (h_B^i + n_{Bw}^i) + \hat{\pi}^V n_{Vw}^i) S^i, \\ S^i - T^i &= S^{i'}, \\ (1 - \pi_d^i - \pi_r) I^i + T^i &= I^{i'}, \\ \pi_r I^i + R^i &= R^{i'}, \\ \pi_d^i &= H(I^y + I^o), \end{aligned}$$

the infection rates,

$$\begin{aligned} \hat{\pi}^B &= \pi^B \frac{\sum_i I^i (h_B^i + n_{Bw}^i)}{\sum_i (S^i + I^i + R^i) (h_B^i + n_{Bw}^i)}, \\ \hat{\pi}^V &= \pi^V \frac{\sum_i I^i n_{Vw}^i}{\sum_i (S^i + I^i + R^i) n_{Vw}^i}, \end{aligned}$$

and the current death rates (for both young and old),

$$\pi_d^i = H(I^y + I^o).$$

The terminal period value is given by

$$V_{\mathcal{T}}(\text{Epi}) = \frac{1 - \beta^{\mathcal{T}^y - \mathcal{T} + 1}}{1 - \beta} \tilde{u} \cdot (S_{\mathcal{T}}^y + I_{\mathcal{T}}^y + R_{\mathcal{T}}^y) + \frac{1 - \beta^{\mathcal{T}^o - \mathcal{T} + 1}}{1 - \beta} \tilde{u} \cdot (S_{\mathcal{T}}^o + I_{\mathcal{T}}^o + R_{\mathcal{T}}^o)$$

where we set the flow utility after the terminal period, $\tilde{u}$, equal to the flow utility of the young absent the epidemic. This formulation of the maximization problem incorporates the entire life span of the young and the old into the analysis without explicitly modelling birth, ageing, and non-epidemic death.

2.1.1 Competitive equilibrium

The framework just described has a competitive-equilibrium counterpart. To abstract from concerns regarding redistribution of consumption between young and old, we assume that the decision unit in the market allocation is a representative family with both old and young individuals. The family does not know who is susceptible, infected, or recovered but knows how many individuals are susceptible, infected, or recovered of each type.

There are two important externalities that the representative family does not internalize. First, the representative family takes the infection rates as given, i.e., $\hat{\pi}_B$ and $\hat{\pi}_V$ are seen as exogenous. Second, the representative family does not internalize their contribution to potential overloading of the hospital system, i.e., the death rates, π_d^i , are also taken as given.

2.2 Vaccination in an epi-econ framework

In this framework, we study an idealized form of vaccination. Only one dose of the vaccine is assumed to immediately lead to full immunity, and the immunity does not wane.

The value of instantaneous vaccination of Δ young individuals (assuming that vaccination can be targeted towards susceptible individuals) is defined as the difference in value between an initial state, $\text{Epi}_t = [S_t^y, I_t^y, R_t^y, S_t^o, I_t^o, R_t^o]'$, and a new state where Δ young individuals have been lifted from susceptible to recovered, $\text{Epi}'_t = [S_t^y - \Delta, I_t^y, R_t^y + \Delta, S_t^o, I_t^o, R_t^o]'$. The value of the vaccinations is defined as $V_t(\text{Epi}'_t) - V_t(\text{Epi}_t)$. The value of vaccinating the old is defined analogously.

We report the average value of vaccination, defined by

$$\text{Average value of vaccination: } \frac{V_t(\text{Epi}'_t) - V_t(\text{Epi}_t)}{\Delta}.$$

We also report the marginal value of vaccination. Letting $\Delta \rightarrow 0$, we arrive at

$$\text{Marginal value of vaccination: } \frac{\partial V_t}{\partial R_t^i} - \frac{\partial V_t}{\partial S_t^i}.$$

The marginal value of vaccinating a young or an old individual is encoded in the gradient of the value function, $\frac{dV_t}{d\text{Epi}_t} = \left[\frac{\partial V_t}{\partial S_t^y}, \frac{\partial V_t}{\partial I_t^y}, \frac{\partial V_t}{\partial R_t^y}, \frac{\partial V_t}{\partial S_t^o}, \frac{\partial V_t}{\partial I_t^o}, \frac{\partial V_t}{\partial R_t^o} \right]'$. Since our solution method calculates the gradient of the value function, the marginal value is thus readily available without further computations.

Note that in our model, neither the planner nor the family head know *who* was vaccinated, only that a share of the young/old were vaccinated.

3 Calibration

This section gives a very short overview of the parameterization and calibration of the model used for the result section. The general idea of the calibration of the economic parameters is to choose parameters to match pre-pandemic data on time use from the American Time Use Survey (ATUS), BLS data on the size of the workforce in the B and V sector respectively, and long-run macroeconomic facts.

The epidemiological parameters are calibrated to best available information in December 2020, and is in line with the parameters used in other contemporaneous studies to facilitate comparisons of results. For a detailed description of calibration strategy and data, see Boppart, Harmenberg, Hassler, Krusell, and Olsson (2021).

Discount rate A period in the model corresponds to a day. We set the discount factor β such that $\beta^{365} = 0.96$.

Demographics We divide the population into young individuals, aged 15-60 years, and old (60+ years old). The old population is less productive, have fewer remaining years of life, and are more vulnerable to the disease. In the beginning of the epidemic, the young constitute a fraction ϕ_0^y of the population, and the old $\phi_0^o = 1 - \phi_0^y$.

Utility v Individuals get utility from consuming the B and V goods together with spending leisure time in the respective sector, as well as from being alive. The flow

utility for an individual is given by

$$v(c_B, h_B, c_V, h_V) = \log \text{CES}(\tilde{c}_B, \tilde{c}_V; \lambda, \varepsilon) + \underline{u}, \quad (1)$$

$$\tilde{c}_B = \text{CES}(c_B, h_B; \lambda_B, \varepsilon_B), \quad (2)$$

$$\tilde{c}_V = \text{CES}(c_V, h_V; \lambda_V, \varepsilon_V), \quad (3)$$

where the constant-elasticity aggregator $\text{CES}(\bullet)$ is defined by

$$\text{CES}(x_1, x_2; \lambda, \varepsilon) = \left(\lambda x_1^{\frac{\varepsilon-1}{\varepsilon}} + (1-\lambda)x_2^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}}. \quad (4)$$

We set $\underline{u}$ so that the statistical value of a life period equals 8 times period consumption.

The production function F Production in both the B and the V sector is given by a Cobb-Douglas function of CES aggregates,

$$F_j(n_{jh}^y, n_{jw}^y, n_{jh}^o, n_{jw}^o) = k_j^\alpha \tilde{n}_{jh}^\nu \tilde{n}_{jw}^{1-\alpha-\nu}, \quad (5)$$

$$\tilde{n}_{jh} = \text{CES}(\phi^y n_{jh}^y, \phi^o n_{jh}^o; \varphi, \theta), \quad (6)$$

$$\tilde{n}_{jw} = \text{CES}(\phi^y n_{jw}^y, \phi^o n_{jw}^o; \varphi, \theta), \quad (7)$$

where the CES function is given by (4). Given our focus on short-run analysis, the sector-specific capital stocks k_B and k_V are fixed. Young and old labor are close substitutes in the production.

Epidemiological parameters The epidemic spreads when people meet for B leisure or in the B or V workplace. The contagiousness of the epidemic is given by the parameters π_V and π_B which we set so that $R_0 = 2.0$ with pre-pandemic behavior. We set the recovery rate, π_r , to $1/18$, so that the average time between infection and recovery is 18 days.

The mortality risk depends on how many currently infected there are, and is substantially higher for the old. We assume that the mortality risk increases substantially once the hospitals are overcrowded: The current death rate in any time period is given by:

$$\pi_{d,t}^i = H(I_t) = \pi_{d,low}^i + \frac{\pi_{d,high}^i - \pi_{d,low}^i}{1 + e^{-k(I_t - \bar{I})}} \quad (8)$$

with $I_t = I_t^y + I_t^o$.

Summary of calibration Table 1 gives a summary of all parameter values.

Parameter	Description	Value
<i>Preference parameters</i>		
β	Discount factor	$0.96^{1/365}$
λ	Weight on $\tilde{c}_B$	0.25
λ_B	Weight on c_B	0.93
λ_V	Weight on c_V	0.69
ε	Elasticity between $\tilde{c}_B$ and $\tilde{c}_V$	1.0
ε_B	Elasticity between c_B and h_B	0.41
ε_V	Elasticity between c_V and h_V	0.80
$VSTP$	Value of a statistical time period (as multiple of period consumption)	8.0
$\underline{u}$	Intrinsic value of life per period (in utils)	3.5
<i>Technology</i>		
α	Capital share	1/3
ν	Home work labor share	0.202
ε_n	Elasticity of substitution between young and old	10
λ_n	Production weight on young	0.62
k_B/k_V	Relative capital stock	0.25
<i>Demographics</i>		
ϕ_0^y	Fraction young	0.73
$\mathcal{T}^y$	Remaining life time young	$31.6 \cdot 365$
$\mathcal{T}^o$	Remaining life time old	$9.2 \cdot 365$
<i>Epidemic variables</i>		
R_0	Spread factor standard SIR model	2.0
$\kappa_B = \kappa_V$	Spread factor economic model	0.24
π_r	Recovery rate	1/18
$\pi_{d,low}^i$	Death rate (before overcrowding) [young, old]	$[0.001, 0.025] \cdot 1/18$
$\pi_{d,high}^i$	Death rate (when overcrowded) [young, old]	$[0.002, 0.050] \cdot 1/18$
<i>Health care system</i>		
ι_h	Fraction of infected in need of hospitalization	0.03
ι_i	Fraction of hospitalized in need of ICU	0.29
ι_b	Inhabitants per ICU bed	3,400
$\hat{I}$	Midpoint logistic function (fraction infected)	$1 / (\iota_h \cdot \iota_i \cdot \iota_b)$
k	Steepness parameter	1,000

Table 1: Summary of calibrated parameters.

4 Results

4.1 Marginal vaccination

In this subsection, we report the time path of the marginal value of vaccination both in competitive equilibrium and in the planner's allocation. We report two measures of the marginal value of vaccination.

First, we report the value of *targeted vaccination*, the marginal per-capita value of a vaccine if we can be sure to vaccinate a currently susceptible individual. This value, in utility terms, is simply the difference in the marginal value of a recovered and a susceptible person,

$$\frac{\partial V_t}{\partial R_t^i} - \frac{\partial V_t}{\partial S_t^i}, \quad i \in \{y, o\}.$$

Second, we report the value of *untargeted vaccination*. This is the marginal per-capita value of drawing a person at random from the young/old population to vaccinate, where the vaccination has an effect if and only if the individual is still susceptible. The value of untargeted vaccination is

$$s_t^i \left(\frac{\partial V_t}{\partial R_t^i} - \frac{\partial V_t}{\partial S_t^i} \right), \quad i \in \{y, o\},$$

with s_t^i denoting the fraction of young/old currently susceptible.

The two measures serve complementary purposes. Since the social planner and the representative family cannot target other behavior and policies conditional on individuals' epidemiological state, it is natural to evaluate the value of untargeted vaccination since it respects the information constraints of the model. However, the value of untargeted vaccination can be low either because the value of vaccinating a susceptible person is low or because it is difficult to find and vaccinate the susceptible person. The results for targeted vaccination show the pure value of vaccinating a susceptible person.

4.1.1 Planner allocation

We show the time path of the marginal value of vaccination under the social-planner allocation in Figure 1. In Figure 1a, we show the evolution of the epidemic under the planner's allocation. The social planner manages the epidemic so that the health system is not overwhelmed while reaching herd immunity without overshooting, i.e., without making the number of recovered and vaccinated individuals exceed what is necessary to reach herd immunity. In Figure 1b, we show the value of a marginal




Figure 1: The marginal value of vaccination under the social-planner allocation.

vaccination for the social planner. We report the value of a marginal vaccination in consumption equivalents, with “days of consumption” as unit. The solid curves show the marginal value of untargeted vaccination while the dashed curves show the marginal value of targeted vaccination.

The marginal value of targeted vaccination is always greater than the marginal value of untargeted vaccination (since the share of susceptible is less than 1). However, the gap between the values of targeted and untargeted vaccination is relatively small for the old. Since relatively few old ever get infected, the susceptible share of old remains high throughout the development of the epidemic. By contrast, the gap between the values of targeted and untargeted vaccination becomes large for the young as the epidemic progresses. Initially, essentially all young are susceptible so untargeted and targeted vaccination are equally effective. The population of young gradually become infected and recovered which reduces the value of untargeted vaccination. Eventually, as the population reaches herd immunity, the value of vaccinations approach zero (there is no virus left to vaccinate against).

The planner prefers to vaccinate the young first, with the value of vaccinating a young individual exceeding 250 days of consumption in the beginning of the epidemic. The value of vaccinating an old individual is less than 200 days of consumption in the beginning of the epidemic. Although the old are more vulnerable, the social planner internalize that the young are more likely to spread the virus.

Toward the end of the epidemic, e.g., day 350, it is more valuable to vaccinate an old first (assuming that vaccination cannot be targeted toward the susceptible). Since




Figure 2: The representative family’s marginal value of vaccination under competitive equilibrium.

the old are still susceptible while most of the young are either infected or recovered, vaccinating the old is simply a de facto more targeted policy than vaccinating the young.

4.1.2 Competitive-equilibrium allocation

Figure 2 shows the time path of the value of a marginal vaccination in competitive equilibrium. In Figure 2a, we show the evolution of the epidemic in competitive equilibrium. The health system becomes overwhelmed and the epidemic is over after one year. During the height of the epidemic, the old have very little in-public leisure and do not work in the workplace. By contrast, the young do change their behavior during the epidemic but to a much smaller degree. Consequently, the spread of the epidemic almost entirely goes through the young population and herd immunity is reached almost entirely through the young becoming infected.

In this competitive equilibrium, the family valuation of the marginal value of a vaccine is shown in Figure 2b. We report the family’s valuation of the marginal value of a vaccine since it is the family’s value function (and not the social planner) which is used to solve the model, and for which we obtain the gradient in the process of solving for equilibrium.

The family values vaccines substantially less than the social planner, at around 80 days of consumption per vaccine in the beginning of the epidemic. Furthermore, the family prefers vaccinating an old individual to vaccinating a young individual. Both targeted and untargeted vaccination of the old are more valuable than vaccinating a

	Young	Old
Family valuation	71 \$7,800	87 \$9,600
Societal value	162 \$17,800	105 \$11,600

Table 2: Marginal value of a vaccination at time $t = 0$ under the competitive equilibrium allocation. Valuation in days of consumption and translated to 2019 USD using annual consumption per capita of USD 40,000 (figure from Q4 2019).

young individual for the entire duration of the epidemic. In practice, since so few old become infected, the difference between targeted and untargeted vaccination of the old is small.

The preference for vaccinating the old reflects that the representative family does not internalize the indirect value of mitigating the spread of the virus through vaccination. Since the young are the primary spreaders in competitive equilibrium, this externality is particularly large for the young. By contrast, the old are more vulnerable if becoming infected, something that the representative family does internalize.

4.1.3 The societal value of a marginal vaccination in competitive equilibrium

From our solution method, the societal value of a marginal vaccination under a social planner and the family’s valuation of a marginal vaccination under competitive equilibrium were both readily available. However, the societal value of a marginal vaccination in competitive equilibrium cannot be backed out of the solution since only the gradient of the representative family’s value function is part of the solution method for competitive equilibrium. Here, we instead compute the societal value of a marginal vaccination at time $t = 0$ in competitive equilibrium by computing a numerical derivative through perturbing the initial conditions, recomputing the competitive equilibrium, and computing total utility.

In Table 2, we display the family’s $t = 0$ valuation of a marginal vaccine together with the societal value of a $t = 0$ marginal vaccine. We also convert the consumption equivalents to US dollars using an annual consumption per capita of \$40,000. The societal value in competitive equilibrium of a vaccine is substantially higher than the family’s valuation, since it incorporates the value of mitigating the spread to others, not only the individual risks associated with becoming infected. While the family values the

vaccination of a young individual to \$7,800, the value for society is \$17,800, more than twice as high. For the old, the societal value of vaccinating is higher (\$11,600) than the family's valuation (\$9,600) but not substantially so since the old in competitive equilibrium are not the primary spreaders of the virus.

From these results, there is a potential tension in implementing a vaccine policy. While society would benefit from vaccinating the young first, any given family would prefer to vaccinate their old first.

4.2 Non-marginal vaccination

In the previous subsection, we studied the value of a marginal vaccination. Although the value of a marginal vaccination is informative for some policies, many vaccination programs are distinctly non-marginal. This raises several questions. For example, how much more valuable is it to distribute the vaccine quickly (in the limit, instantaneously) rather than gradually? Does it matter how the doses are distributed (e.g., to the young or the old) and does this prioritization depend on the roll-out speed? What is the social cost of a subset of the population who refuse to take a vaccine? Our framework is well suited to answer these types of questions. Here, we consider a specific thought experiment about mass vaccination at time $t = 0$.

We compute the societal value in competitive equilibrium of vaccinating a share of the population at time $t = 0$ with the vaccinated group either all young or all old. Computationally, it is straightforward to analyze any time period t during the epidemic, to look at the planner's allocation, and to consider vaccinating both young and old at the same time. We calculate the average value of vaccinations. For example, the average value of vaccination Δ young individuals is given by

$$\begin{aligned} \text{Epi}_t &= [S_t^y, I_t^y, R_t^y, S_t^o, I_t^o, R_t^o]', \\ \text{Epi}'_t &= [S_t^y - \Delta, I_t^y, R_t^y + \Delta, S_t^o, I_t^o, R_t^o]', \\ \text{Average value of vaccination} &= \frac{V_t(\text{Epi}'_0) - V_t(\text{Epi}_0)}{\Delta}. \end{aligned}$$

In Figure 3, we display the value of vaccinating a share of the population. Near $\Delta = 0$, the average value of vaccinating is increasing in the number of vaccinated, a form of increasing returns to scale. Vaccinations reduce both the overshooting of the epidemic and the period when the hospital system is overwhelmed. In a large neighborhood of $\Delta = 0.0$, vaccinating the young gives a higher value than vaccinating the old. However, vaccinating all the old, or almost all the old, gives a higher value




Figure 3: The per-vaccination average value of vaccinating a share of the population, either all young or all old, at time $t = 0$.

than vaccinating a corresponding number of young individuals.

4.2.1 Vaccinating 20% of the population, all young

If 20 percent of the population, all young, are vaccinated, much of the spread is halted and more old individuals dare to go out. In Figure 4, we show the evolution of the epidemic under such a scenario. The hospital system is only overwhelmed for a short period and the average value of the vaccines is large, 225 days of consumption per vaccination. Where are these benefits coming from? In our model with endogenous behavioral responses, a substantial part of the value of vaccination is that people enjoy more in-puBlic leisure when the risk of infection is lower. 40 percent of the total benefit of vaccinating the young stems from increased flow utility for the old, who enjoy more in-puBlic leisure. The remainder, 60 percent, of the benefits of vaccinating the young stems from a lower death rate for the young (a reduction from 0.14 percent to 0.06 percent). This effect comes both from the direct effect of the vaccination regime, the less overwhelmed hospital system, and the reduced overshooting. Perhaps surprisingly, the death rate among the old is slightly higher (0.46 percent instead of 0.44 percent) when 20 percent of the population, all young, are vaccinated. When the risk of infection falls, the old enjoy more in-puBlic leisure and in our calibrated model this behavioral response is so strong that the net effect is that the mortality rate for the old slightly




Figure 4: The average value of vaccinating 20% of the population, all young, and the infection curve associated with such an intervention.

increases.

4.2.2 Vaccinating 22% of the population, all young

The maximum average value of vaccines for young is obtained when 22% of the population, all young, are vaccinated. This maximum is obtained when the hospital system is no longer overcrowded, as can be seen in Figure 5.

4.2.3 Vaccinating 25% of the population, all old

Vaccinating 25 percent of the population, all of them old, yields very large benefits, as can be seen in Figure 6. Approximately 60 percent of these benefits come from the higher flow utility of the old. As the old are mostly vaccinated, they no longer need to shield at home to avoid becoming infected. As a result, they are able to spend more time on in-public leisure with substantial increase in utility. The remaining 40 percent of the benefits are evenly split between a lower death rate among the old and a lower death rate among the young. The death rate for the old is reduced from 0.44 percent (without a vaccine) to 0.14 percent. The death rate for the young is reduced from 0.14 percent to 0.09 percent. The young have a lower death rate because (i) fewer individuals need to become infected to reach herd immunity, (ii) there is less overshooting the herd immunity threshold, and (iii) the hospital system is less overcrowded. The hospital system is still overcrowded for a period of time, in contrast




Figure 5: The average value of vaccinating 22% of the population, all young, and the infection curve associated with such an intervention.

to the previous scenario vaccinating young individuals constituting 22 percent of the population. Because the young are the primary spreaders, vaccinating the old have a comparatively smaller effect on slowing down the epidemic.


4.2.4 Vaccinating 27% of the population, all old

Vaccinating all the old individuals ensures that no old individuals become infected. Still, the epidemic does affect the old negatively as well. In this scenario, the family planner sends the old to work more and their leisure time decreases. During the epidemic, the flow utility of the vaccinated old falls more than the flow utility of the young.


In our calibration, vaccination of all the old happens to ensure that the hospital system is not overcrowded, as can be seen in Figure 7. As a result, vaccinating the old brings down the death rate of the young from 0.14 percent (without any vaccination) to 0.06 percent.

4.2.5 Vaccinating 45% of the population, all young

Finally, if 45 percent of the population, all young, are vaccinated, the epidemic never takes off. The population has reached herd immunity already at time 0. After this point, the average value of vaccines is decreasing quickly since the marginal value of an extra vaccination is zero. In this model framework, there is no reason to vaccinate




(a) Average value of vaccines




(b) Infection curve if 25% old are vaccinated at time $t = 0$

Figure 6: The average value of vaccinating 25% of the population, all old, and the infection curve associated with such an intervention.



(a) Average value of vaccines



(b) Infection curve if 27% (old) are vaccinated at time $t = 0$

Figure 7: The average value of vaccinating 27% of the population, all old, and the infection curve associated with such an intervention.




Figure 8: The average value of vaccinating 45% of the population, all young, and the infection curve associated with such an intervention.

additional individuals once it is ensured that the epidemic will not take off. The value of vaccinating enough people at time $t = 0$ so that the epidemic never takes off is approximately 28% of annual total consumption ($\frac{0.45 \times 225}{365} = 28\%$).

5 Concluding remarks

We have introduced a framework for integrated epi-econ assessment and shown how the framework can be used to analyze vaccination policy. In particular, we have used the framework for comparing the prioritization order of groups for the vaccine distribution. Our results underline that the trade-offs not only involve output vs. lives but, importantly, take into account how social leisure is affected.

Our vaccine evaluations focus on a quick roll out of the vaccine program. In terms of our model, we consider vaccination (of one or more citizens) right at the beginning of the pandemic. In practice, policymakers did not face this question, as vaccines were not available at that time, and rolling out programs took time and faced many challenges. (In this sense, analysis of the precise experiments we conduct in this paper are perhaps best suited for the next epidemic, i.e., one that we are prepared, and have vaccines, for.) Experiments with vaccination at later stages is, however, completely straightforward to carry out with our framework. One can also assume that the health care sector is a constraining factor and needs significant amounts of time to roll out the vaccine program. The quantitative assessment of the dollar-equivalent value of one

vaccination would clearly vary with the timing. It seems highly unlikely, however, that the value would fall so much that their net value (i.e., after deducting production costs) becomes negative. It also seems unlikely that the prioritization issue loses importance, and the preferred order may depend on the timing of the vaccine.

It would be important to extend our work in other directions: one should introduce more features of the complex reality we live in. Let us therefore emphasize a few weaknesses of our present setting. One is that the calibration of the epidemic is not based on the latest developments (e.g., we use $R_0 = 2$ while the delta variant, which occurred later, is much more contagious). With a higher R_0 , the herd immunity threshold increases; with a homogeneous population, a R_0 of 2 implies herd immunity at 50%, while an R_0 of 8 raises it to 87.5%. More generally, it appears that epidemics often evolve in this manner, i.e., through mutations and corresponding new waves. Assumptions about the contagiousness, including how it evolves over time, are fortunately easy to change in our model, and we keep the “pre-delta” calibration here only to facilitate comparison with other studies. Our setting also builds on perfect protection immediately, that immunity does not wane, and that the vaccinated are not contagious; these features make us overstate the value of a vaccine, and it would be interesting—and straightforward—to relax these assumptions. To include anti-vaxxers (by assuming that a fraction of the population “cannot” be vaccinated) would be another, valuable, extension.

Our general framework, of course, has other shortcomings. The age structure, along with age-dependent patterns of social activities, can be made much richer. An important first step would be to include a notion of nursing homes and their particular features—combinations of points of interaction between young and old—that appear to have been central during the present pandemic. A much harder extension, not only for us but all other epi-econ models as well, is to drop our maintained information assumption that individuals do not know their epidemiological state. In a realistic model, the economy-wide state variable should include the full distribution of individual beliefs about themselves (and their beliefs about others...), which would be endogenous and evolve over time; moreover, it would not make sense to assume that a fictitious planner knows this distribution. Clearly, our informational assumption makes the model tractable. Are there steps toward introducing non-trivial beliefs and their updating? In our original paper, we outline some “tricks” that could be used to move in this direction and that could turn out useful. There, we also discuss how proper social networks could be included in a computationally feasible way.

The overarching purpose of our research is to move toward the construction of

realistic-enough assessment models that can be put to productive use in practice, especially by policymakers and health agencies. Thus, the purpose is not to emphasize qualitative features or to discover new mechanisms—though of course our work could generate new insights of this sort too, as a byproduct. As a part of this model development we aim to make the computational tools accessible: plug-and-play, a Dynare of sorts for the epi-econ area. Already at this point, we are quite close to that.

References

- Arellano, C., Y. Bai, and G. Mihalache (2021). Covid-19 vaccination and financial frictions in emerging markets.
- Auld, C. and F. Toxvaerd (2021). The great covid-19 vaccine rollout: Behavioral and policy responses. Technical report, CEPR Discussion Paper No. DP16070.
- Azzimonti, M., A. Fogli, F. Perri, and M. Ponder (2020). Pandemic control in econ-epi networks. Technical report, National Bureau of Economic Research.
- Boppart, T., K. Harmenberg, J. Hassler, P. Krusell, and J. Olsson (2021). Integrated epi-econ assessment: quantitative theory.
- Chen, F. and F. Toxvaerd (2014). The economics of vaccination. *Journal of Theoretical Biology* 363, 105–117.
- Garriga, C., R. Manuelli, and S. Sanghi (2021). Optimal management of an epidemic: Lockdown, vaccine and value of life. Technical report.
- Glover, A., J. Heathcote, D. Krueger, and J.-V. Rios-Rull (2021). Optimal age-based vaccination and economic mitigation policies for the second phase of the covid-19 pandemic.
- Goodkin-Gold, M., M. Kremer, C. M. Snyder, and H. L. Williams (2020). Optimal vaccine subsidies for endemic and epidemic diseases. Technical report, National Bureau of Economic Research.