

Delli Gatti, Domenico; Reissl, Severin; Turco, Enrico

Working Paper

V for Vaccines and Variants

CESifo Working Paper, No. 9291

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Delli Gatti, Domenico; Reissl, Severin; Turco, Enrico (2021) : V for Vaccines and Variants, CESifo Working Paper, No. 9291, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/245472>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

V for Vaccines and Variants

Domenico Delli Gatti, Severin Reissl, Enrico Turco

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

V for Vaccines and Variants

Abstract

We employ a new version of the ABC macro-epidemiological agent based model presented in Delli Gatti and Reissl (2020) to evaluate the effects of vaccinations and variants on the epidemic and macroeconomic outlook. Vaccination plays the role of a mitigating factor, reducing the frequency and the amplitude of contagion waves, while also significantly improving macroeconomic performance. The emergence of a variant, on the other hand, plays the role of an accelerating factor, increasing the volatility of epidemic curves and worsening the macroeconomic outlook. If a more contagious variant emerges after vaccination becomes available, therefore, the mitigating factor of the latter is at least partially offset by the former. A new and improved vaccine in turn can redress the situation. Vaccinations and variants, therefore, can be conceived of as drivers of an intertwined cycle impacting both epidemiological and macroeconomic developments.

JEL-Codes: E210, E220, E240, E270, I120, I150, I180.

Keywords: agent-based models, epidemic, Covid, vaccination, variant.

Domenico Delli Gatti
Department of Economics and Finance &
Complexity Lab in Economics
Catholic University / Milan / Italy
domenico.delligatti@unicatt.it

Severin Reissl
IUSS Pavia / Italy
severin.reissl@gmail.com

Enrico Turco
Department of Economics and Finance &
Complexity Lab in Economics
Catholic University / Milan / Italy
enricomariaturco@gmail.com

August 26, 2021

Special thanks to Giulia Gardin for excellent research assistance.

1 Introduction

Since the fall of 2020, two crucial V-words have changed the epidemiological (and macroeconomic) landscape worldwide: Vaccines and variants. Since the outbreak of Covid 19, vaccination has been considered the ultimate weapon against the virus. The emergence of virus variants has put this belief to the test as mutations may eventually reduce the efficacy of the vaccine. In this paper we explore the consequences of these phenomena using a new version of ABC, the agent based macro-epidemiological model of Covid-19 put forward in Delli Gatti and Reissl (2020).

We proceed in two steps. First, we analyze the effects of vaccination in the absence of any variant. We then introduce and study the effects of variants in the presence of the original vaccine and of an improved vaccine.

From the epidemiological point of view vaccination may affect the spread of the disease along three dimensions. First and foremost, vaccination may reduce the probability for a vaccinated person to get the disease. The lower this probability, the higher the efficacy of the vaccine along this dimension. We term this Vaccine Efficacy of type 1 (VE_1).

Second, the vaccine may reduce the probability for a vaccinated person to develop serious symptoms and eventually die if they are infected. The lower this probability, the higher the efficacy of the vaccine along this dimension, which we term Vaccine Efficacy of type 2 (VE_2).

Third, the vaccine may reduce the probability for a vaccinated but infected person to infect a susceptible individual. The lower this probability, the higher the efficacy of the vaccine along this dimension, which we term Vaccine Efficacy of type 3 (VE_3).

Ample evidence shows that that VE_1 for vaccines against Covid 19 is not complete; estimates of the induced immunity to infection vary from 60 to 90%, also depending on the vaccine type (smaller for traditional vaccines, higher for mRNA vaccines). It is also widely observed that VE_2 is very high (in the neighborhood of 90%): in the summer of 2021, in fact, the vast majority of people hospitalized and dying from the disease were unvaccinated. Finally, there is some evidence that vaccinated but infected individuals can still infect susceptible individuals but estimates still differ widely.¹

We expect the epidemiological consequence of vaccination to be a decline in (a) the incidence/morbidity and (b) the mortality of the disease. The magnitude of these effects may of course also depend on the particular configuration of the vaccination campaign, and we experiment with different prioritization rules in order to investigate this.

As the macroeconomic dynamics of our model are strongly affected by epidemiological developments, we expect vaccination (1) to reduce the duration and amplitude of the output loss due to the epidemic; (2) to alleviate the burden of public sector deficit and

¹In our simulations we have assumed that vaccine efficacy of type 1 and 2 are $VE_1 = 70\%$ and $VE_2 = 90\%$ respectively. As to vaccine efficacy of type 3 we put ourselves in the worst case scenario and suppose that it is zero: a vaccinated infected individual is as infective as an unvaccinated one.

debt.

We obtain a rich set of results from the simulation of the vaccination campaign. First of all we infer that the vaccine significantly contributes to containing the epidemic and saving lives, in line with expected results (a) and (b) above. Regardless of the prioritisation strategy, in fact, both the cumulative number of infections and of deaths is substantially lower than in the absence of a vaccine.

Second, simulations show that in the presence of a vaccine, while there are additional waves of infections (due also to the fact that vaccine-induced immunity is temporary), both the frequency and amplitude of these waves are reduced. Moreover, in the presence of the vaccine the curve of new deaths remains mostly flat while there are multiple endogenous waves of newly deceased in the absence of a vaccine. Vaccination is hence a significant *mitigating factor* of the cyclical dynamics of infections and deaths, leading to a reduction of the frequency and the amplitude of subsequent waves once the vaccine becomes available.

Third, the epidemiological results of different prioritization strategies do not differ as starkly as one might expect: in our simulations, the magnitude of the mitigation effect of vaccination is similarly strong across different prioritization strategies. Priority given to the old, however, allows to save a somewhat higher number of lives at the cost of a slightly higher number of infections. A vaccination strategy aimed at prioritizing the old protects individuals with higher exposure to the risk of dying while younger agents are more exposed to infection due to interactions at the workplace.

Fourth, the vaccine has a relatively small but significant and persistent positive effect on GDP, particularly since the lower number of infections reduces the (negative) consumption shock associated with social distancing.

We then turn to virus variants. A variant may change the epidemiological scenario along four dimensions: it may (i) increase the probability for a susceptible individual to be infected if in contact with an infectious individual; (ii) increase the probability for an infected individual to develop serious symptoms and eventually die; (iii) reduce the duration of the immunity acquired by people who have recovered from the disease; (iv) reduce the efficacy of the vaccine along one or more of the dimensions outlined above.

For simplicity, we assume that variants alter the epidemiological scenario only along dimensions (i) and (iv). In particular, we experiment with two variants. Both of them are more contagious than the original virus. In addition, both variants reduce the efficacy of the original vaccine in preventing infection (VE_1). The second variant in addition also reduces the efficacy of the original vaccine in preventing serious symptoms and death (VE_2). We explore the effects of both variants both in the absence and presence of an *improved* vaccine which eliminates these latter effects of the variants.

Simulations show that in all cases, the variant replaces the original virus very soon, driving a new wave of infections. This increase in infections also leads to an increase in deaths

even in the case of variant 1 which does not affect the vaccine’s ability to prevent serious disease. As one would expect, however, the death toll increases even more strongly in the case of variant 2. Variants hence act as an *accelerating factor*, leading to an increase of the frequency and amplitude of waves of contagion, partly offsetting the positive effects of the original vaccine. If the vaccine is adapted to the variant, the amplitude and frequency of these waves is once more mitigated.

The paper is structured as follows. After a brief review of the literature (section 2), we present a synthetic overview of the ABC model in section 3. We provide a detailed description of the model in appendices A and B. Section 4 discusses the calibration and validation of the baseline scenario. Section 5 presents the epidemiological scenarios while section 6 focuses on the macroeconomic effects of these scenarios. We discuss the effects of vaccination in section 7. Section 8 introduces the emergence and spread of variants. Section 9 concludes.

2 Related literature

Even before the outbreak of Covid 19, many papers highlighted the economic benefits of vaccination. Smith et al. (2009) use a computable general equilibrium model to show that the availability of an effective vaccine is key in determining the economic impact of influenza in the UK. Similarly, Keogh-Brown et al. (2010) use a multi-sector computable general equilibrium model to examine the potential economic cost of a modern epidemic. They find some evidence that vaccinations, antivirals or a mix of the two, may be a cost-effective strategy.

Since the outbreak of the current pandemic, some studies have specifically analyzed the macroeconomic effects of a vaccine against Covid 19. Using a SIR model to address the question of prioritization in vaccination, Forslid and Herzing (2021) find a trade-off between economic gain and fatality reduction. Economic gain is maximised when the middle-aged are given priority. When instead the elderly prioritised, deaths are minimised. Policy choices therefore depend on priorities: reduction of fatalities or maximisation of the economic gain. In a different framework, Babus et al. (2021) reach similar conclusions. They suggest a vaccine distribution that puts more emphasis on the risk of mortality than on the loss of economic output. Gollier (2021) also use a SIR model, concluding that it is key for the number of lives saved and for the economy to vaccinate older people first. According to Matrajt et al. (2021), prioritization depends on vaccine efficacy. In order to minimize deaths, when vaccine efficacy is “low”, it is optimal to allocate vaccines to the old first. On the contrary, when vaccine efficacy is high, priority should be given to younger age groups.

Lindskog and Strid (2020) assess the economic gains of a rapid vaccination programme against a benchmark consisting of the economic costs of the epidemic in terms of loss of

production, higher public debt and decline in employment that the Swedish economy has faced in 2020. They identify as most important the indirect costs deriving from lockdowns and social distancing measures. They focus on the effects that rapid vaccination has on central macroeconomic variables and find that shortening the duration of the epidemic through vaccination would allow to recover major economic losses.

Finally, Saad-Roy and Wagner (2020) find that contagion dynamics can be altered substantially by a mass roll-out of vaccines in a SIR model. The impact, however, is strongly dependent on the efficacy of the vaccine and the response of the immune system. But even with imperfect immunity and moderate vaccination rates, vaccination could be key in strengthening the control of the epidemic.

With this paper we contribute to a small but growing literature on joint economic and epidemiological dynamics in agent-based settings (e.g. Mellacher, 2020; Basurto et al., 2021). The model presented by Basurto et al. (2021) includes the possibility of vaccination, but the authors do not conduct experiments to discern the specific impact of vaccines on economic performance or epidemiological dynamics as it is assumed that all agents are immediately vaccinated once the vaccine becomes available, and that the latter is fully effective. They do however show that the emergence of a more contagious variant may have substantial impacts both in terms of epidemiological and economic outcomes.

3 An overview of the model

3.1 The environment

The model economy we analyze is populated by households, firms, the banking system and the public sector. The unit of time for the macroeconomic sub-model is a month. The epidemiological sub-model instead runs at the basic frequency of one week, with every month in the model containing exactly four weeks. In the macroeconomic submodel, the time subscript t indicates a month while in the epidemiological sub-model the time subscript τ indicates a week.

There are N_H households which fall into two categories: N_W workers and N_F firm owners. There is a single owner per firm. For simplicity we assume that only workers can become ill, both in “normal times” and during an epidemic. During “normal times” workers may become ill but this illness is neither lethal nor can it be transmitted to others, such that all who become ill eventually recover and the population of living workers will not change. During the epidemic the population of *living* workers may change since the epidemic disease is potentially lethal. If some worker $h \in N_W$ dies they essentially become dormant, and all their assets are written off. Over time, dead agents are replaced by newly born ones which take their place, such that the population of living workers may once more reach N_W .

There are N_F firms which fall into three categories: N_F^k producers of capital goods (K-firms), N_F^b producers of *basic* (or essential) consumption goods (B-firms) and N_F^l producers of non essential consumption goods or *luxury goods* (L-firms). In the following we will consider also the set of all consumption goods producers (C-firms) which is the union of the sets of B-firms and L-firms, denoted N_F^c . The number of *active* firms may change over time due to entry and exit. A firm exits if it goes bankrupt, i.e., if its equity turns negative. Over time, bankrupt firms may be replaced by new entrants in the same sector, but the total number of firms never exceeds N_F .

The banking system is represented by a single bank, collectively owned by firm owners. In principle the bank can go bankrupt. If the bank's equity turns negative we assume that its owners will provide an equity injection in order to make it survive.

3.2 The macroeconomic sub-model

In this section we succinctly describe the macroeconomic sub-model. We present behavioural rules for the four groups of economic agents (households, firms, the banking system and the public sector) and market protocols that govern transactions. A more detailed description is given in appendix A.

3.2.1 Households

A household indexed with $h \in (1, N_W)$ is a worker. If alive, workers can be either economically active or economically inactive. Chiefly for epidemiological purposes, the population of workers is divided into three age groups: young, middle-aged and old. All old agents are assumed to be retired and hence economically inactive. All young and middle-aged agents are initially economically active and constitute the labour force (either employed or unemployed). When an economically active worker falls ill they become economically inactive until their illness ends.²

Each economically active worker supplies 1 unit of labour inelastically. If employed, they receive a uniform wage and pay a fraction (the tax rate on wages) of this wage to the Government. If unemployed, they will receive an unemployment subsidy. Workers (both employed and unemployed) who fall ill receive sick-pay. Retired workers receive pensions. A household indexed with $h = N_W + f$ is the owner of the f -th firm, $f = 1, 2, \dots, N_F$. The income of this household consists of dividends, which, in turn, are equal to a fraction (the pay-out ratio) of the after-tax profit of the firm owned by that household. The firm pays out dividends only if profits are positive. Moreover, the firm owners are assumed to jointly own the bank and consequently each one of them receives an equal share of the dividends distributed by the bank.

²In normal times, illness always ends with recovery, during an epidemic it may end either with recovery or with death.

In addition, all households receive interest income on any deposits held at the bank, which represent the only financial assets owned by households.

A household's *consumption budget*, i.e., the sum it wishes to spend on consumption goods, is given by a weighted average of past disposable incomes and a fraction of its financial wealth (deposits). The fraction of the consumption budget allocated to B-goods depends on the relative price of B-goods and L-goods. The consumer shops first at B-firms and then at L-firms.

The consumer visits two B-firms: the “go-to” supplier and a randomly drawn potential new shopping partner. If the price charged by the former is lower than or equal to that of the latter they will first buy from the go-to supplier and resort to the new seller only if the consumption budget devoted to B-goods is not completely exhausted with the first purchase. Otherwise, they will switch to the new partner (and reverse the order of purchase) with a probability which is increasing with the price set by the go-to partner relative to that of the potential new partner. If the consumer actually switches to the new partner, the latter becomes their new go-to supplier. If a firm goes bankrupt and exits (see below), all the households who purchased at this firm as their go-to supplier will randomly choose a different one. The market protocol for L-goods follows the same rules as that for B-goods.

This market protocol does not guarantee equilibrium. Queues of unsatisfied consumers (involuntary savers) at some firms may coexist with unsold goods at some other firms.

3.2.2 Firms

B-firms and L-firms are consumption goods producers (C-firms for short) and follow the same behavioural rules. An active firm indexed with $f \in (1, N_F^c)$ has *market power* and sets the individual price and quantity under uncertainty..

The decision process is based on two *rules of thumb* which govern price changes and quantity changes respectively. Excess demand and the relative price $\frac{P_{f,t}}{P_t}$ – where P_t is an aggregator of the prices set by C-firms – dictate the *direction* of price adjustment: the firm will increase (reduce) the price next period if it has registered excess demand (supply) and has underpriced (overpriced) the good in the current period. Otherwise it will leave the price unchanged. The *magnitude* of price adjustments is stochastic. We also assume that the firm will never set a price lower than its average cost.

Both the direction and the magnitude of quantity adjustment are determined by excess demand. The firm will increase production next period if it has registered excess demand (in the form of a fringe of unsatisfied consumers) in the current period; it will downsize production if it has registered excess supply.

Technology is represented by a Leontief production function the arguments of which are capital and labour. Once a decision has been taken on desired output, the firm retrieves

from the production function how much capital and labour it needs to reach that level of activity. If actual capital is greater than the capital requirement, the desired *rate of capacity utilization* will be smaller than one. If actual employment is smaller than labor requirement, the firm will post vacancies. If the opposite holds true the firm will fire workers. In this scenario, provided there are no bottlenecks on the labour market, the firm can reach the desired level of production.

On the other hand, if actual capital is smaller than the capital requirement, the former will be utilized at full capacity (the rate of capacity utilization will be equal to 1) but desired output will not be reached and scaled back accordingly.

Actual capital in the current period is determined by capital investment carried out in the previous period and the firms' undepreciated capital stock. We assume that a C-firm may carry out investment in any given period with a probability $\pi^k < 1$. In order to determine its investment demand, the firm calculates a target capital stock based on past utilisation and the target utilisation rate, also taking into account the depreciation of capital and the probability of investing. It invests in capital goods so as to reach this target capital stock.

Once investment demand has been determined, C-firms visit the market for K-goods. The market protocol for this market follows the same rules as those for B-goods and L-goods.

The price adjustment rule for capital goods producers is the same as that of C-firms. The quantity adjustment rule departs from the one adopted by C-firms to take into account the assumption that K-goods are durable and therefore storable: inventories of capital goods can be carried on from one period to another and sold in the future. We assume that K-firms are endowed with a linear production function with labour as the only input. Once the price-quantity configuration has been set, a K-firm may post vacancies or fire workers in order to fulfil labour requirements.

The market for labour is characterized by *search and matching*. Unemployed workers visit a subset of firms chosen at random. Once an unemployed worker finds a firm with an unfilled vacancy they stop searching and the match occurs. The uniform nominal wage is set on the basis of labour market conditions captured by the distance between the current unemployment rate and a threshold unemployment rate. Whenever the unemployment rate is above (below) the threshold the wage will decrease (increase).

3.2.3 The banking system

Firms register a financing gap when outlays (to pay for wages and capital investment) are greater than their available liquidity in the form of accumulated bank deposits. Firms which cannot self-finance their costs demand bank loans.

The bank sets the interest rate on loans and the quantity of credit supplied to firms. The

interest rate on loans is set adding a mark up (*external finance premium*) on the risk free interest rate. The external finance premium, in turn, is increasing with the borrower's leverage. Moreover, the bank determines a maximum amount it is willing to lend to a given borrower, again based on that borrower's leverage. This means that a firm may be credit rationed and therefore forced to scale down production and/or investment. In every period, borrowers repay a fraction of their outstanding loans.

Households and firms hold deposits at the bank. The interest rate on deposits is a fraction of the fixed risk free interest rate which coincides with the interest rate on Government bonds. If the bank's profit at the end of a period is positive, it pays a fraction of its after-tax profit as a dividend, which is divided up equally among all firm owners.

3.2.4 The public sector

The public sector collects taxes on wage income and profits and provides transfers in the form of unemployment subsidies, sick-pay and pensions, all of which are given by fractions of the current nominal wage. Government expenditure consists of public provision of healthcare services. In case of a public sector deficit, the Government issues bonds. It is assumed that the bank purchases all issued bonds at a fixed interest rate.

3.2.5 Demand and supply of healthcare

Real government expenditure on healthcare is given by a constant fraction of full employment output, calculated using the initial labour force. We assume that the government uses this amount to spend on the output of both K-firms and C-firms. This expenditure may be interpreted as purchases of equipment necessary for healthcare provision and consumption of medical workers (who are not explicitly modelled). The goods thus purchased are converted one-for-one into a supply of healthcare.

We assume that even in the absence of an epidemic, a healthy worker may become ill with some probability in any period, but that such illness is neither potentially lethal nor infectious to others. As long as an agent is ill, they generate a demand for healthcare which is increasing with their age. They enter the healthcare system to receive treatment. If the remaining supply of healthcare in a given period is insufficient to accommodate the agent's demand, they join a randomised queue to receive treatment.

In the case of an epidemic, agents who contract the epidemic disease and develop serious symptoms will also demand healthcare, making it more likely that demand will exceed supply.

3.2.6 Entry and exit

While during 'normal times' the population of living workers is constant, the epidemic disease may lead to the death of workers. If a worker dies, their assets are written off, and

they may be replaced in each future period by a young worker with a constant probability. During the epidemic, the flow of new deaths may hence at least for some time exceed the flow of new ‘births’, meaning that the population of living workers is smaller than N_W . If a firm’s equity becomes negative, it is assumed to go bankrupt and exit.³ In every future period, the exiting firm may then be replaced by a new firm operating in the same sector, with a probability that is increasing in the average profit rate in the sector in question. The new firm receives any fixed capital remaining from the bankrupt firm and receives an injection of liquidity from the owner of the bankrupt firm, who becomes the owner of the new firm.

The bank’s equity may become negative due to persistent loan defaults. If this is the case, a bail-in procedure is applied: all firm-owners (who collectively own the bank) make a transfer to the bank until its equity becomes positive.

3.3 The Epidemiological sub-model

In this section we briefly describe the epidemiological sub-model. We present a taxonomy of epidemiological segments of the population, the transmission of the disease and the possible outcomes for each infected agent and for the composition of the population. A more detailed outline can be found in appendix B.

3.3.1 The taxonomy of epidemiological groups

The epidemic is characterised by the outbreak of an *infectious* disease which spreads from one subject to the others through *contagion*. At a certain point during the model simulation, a small number of workers are exogenously infected with the epidemic disease. These people are the *initial infected* (and infectious). The susceptible agents after the appearance of the infected are simply the difference between the initial population and the initial infected.

Infected agents can be either non-symptomatic or symptomatic. The former are infected agents who do not have symptoms or develop only mild symptoms. In this case the infection can be detected only if the agent is subjected to a test. The probability of detecting an infected non-symptomatic agent is exogenous in the scenario of *Uncontained Epidemic (UE)* and endogenous in the scenarios of (spontaneous) *Social Distancing (SD)* and *Lock-down (LD)*, becoming a function of the number of newly detected cases in the past week. The scenarios are described in detail in section 5. Detected non-symptomatic infected agents are quarantined and therefore cannot spread the disease. Non-symptomatic and

³If a firm’s liquidity (bank deposits) is negative but its equity is positive, it receives a transfer from the firm owner up to the financial wealth of the owner. If the firm’s liquidity is then still negative, the bank takes a loss equal to the negative balance and the firm’s deposit become zero, but the firm does not exit.

un-detected infected people can still spread the disease. For simplicity, we assume that the infected remain contagious for the entire duration of the illness.

By assumption, all symptomatic infected agents develop serious symptoms and are detected with certainty. The probability to develop serious symptoms is increasing with age. All agents whose infection is detected (both non-symptomatic and seriously ill) will be inactive (and receive sick pay if they are not retired) and will not have social contacts for the entire duration of the disease. Only people developing serious symptoms express demand for healthcare services. For simplicity, we assume that people who are quarantined at home do not need healthcare. Non-symptomatic infected agents recover with certainty after a certain number of weeks. The only agents who may die from the disease are those who develop serious symptoms.

3.3.2 Contagion

We adopt a granular approach to contagion focusing on networks in order to depict the transmission of the epidemic among agents. Contagion takes place in three networks: the workplace (employment network), the marketplace (shopping network) and social relations (social network). Employed workers are nodes in the *employment network*. Each employed workers is linked to all co-workers in the firm they work for. If a worker is infective, they can spread the disease to their (susceptible) co-workers. If a firm goes into smart-working (see below), a fraction of all the possible workplace connections at that firm will be eliminated.

In addition, all workers are nodes in the *shopping network*. A certain number of households shop at a given C-firm. If one of these buyers is infective, they may spread the disease to other households shopping at the same firm. We list all possible connections between the customers of a given firm and assume that a fixed share of those encounters actually take place (not all customers visit the firm at the same time). This share is reduced if there is a lockdown in place or people engage in social distancing.

Contagion also occurs during leisure time. To capture this process we build a *social network*. Each worker has a set of social connections consisting of family and close friends. The total number of social connections is a (very small) fraction of the maximum number of possible undirected connections between workers.

We assume that each infected agent meets all the agents they are connected to (at work, while shopping and during leisure time) in every week. Given the number of connections in week τ which involves exactly one infected and one susceptible agent, we assume that only a fraction (the transmission rate) of these connections may lead to a new infection. This set of *potential* new infections is constructed by randomly drawing from the set of all connections involving one infected and one susceptible agent. We assume that the different types of connections have different probabilities of being drawn, being highest

for social connections, second highest for workplace connections and lowest for market connections.

Each of these *potential* new infections leads to an actual new infection with a probability which changes with the epidemiological scenario: it is one in the scenario of Uncontained Epidemic (UE) and potentially smaller than one in the scenarios of Social Distancing (SD) and Lockdown (LD).

3.3.3 Recovery, death, re-infection

For each infected agent, the duration of the disease is governed by a random variable with finite support. Non-symptomatic agents recover with certainty after this duration. Agents with serious symptoms may die in each period in which they are infected, with a probability which increases both with age and their excess demand for healthcare.

When the healthcare system becomes overburdened, the demand for healthcare is rationed and an agent who develops serious symptoms may be forced to join a randomised queue of ill agents who have an excess individual demand for healthcare. If an agent with serious symptoms has not died after the duration of the disease, they will recover. We assume that both the effect of age and of excess demand for healthcare on the probability of dying decrease over time as the healthcare system is partly able to adapt to dealing with the novel disease.

Recovered agents who became economically inactive due to the disease will re-enter the labour force as unemployed workers and look for a job. The recovered may become susceptible to the epidemic disease again in each future period with a low probability.

4 Calibration and baseline simulation

We construct a *baseline scenario* of the model which corresponds to **Normal Times**, i.e., a situation in which there is no epidemic but people may become ill without dying or being able to infect others. In order to calibrate the parameters for the macroeconomic sub-model, we make use of data for real GDP, consumption, gross fixed capital formation and employment for the Lombardy region from 1995 to 2017, available from the website of the Istituto Nazionale di Statistica (Istat).⁴

At the regional level, data for GDP and its components are available only at annual frequency. We use moments and statistics calculated from these data in order to calibrate and validate the model. We apply the HP filter ($\lambda = 100$) to the empirical time series and then calculate the standard deviations (relative to the trend component) and auto-correlations of the filtered series. Table 1 shows the empirical statistics we obtain as well

⁴We follow the same calibration methodology used for the original ABC model presented in Delli Gatti and Reissl (2020). Since the current model differs from ABC both in terms of the number of agents and some behavioural assumptions, the results of this calibration procedure are also somewhat different.

as 95% confidence intervals (in parentheses) which we generate using bootstrapping.

Table 1: Empirical moments & statistics for Lombardy (1995-2017)

Statistic	GDP	Consumption	Investment	Employment rate
Std. deviation	0.015722 (0.01566; 0.01578)	0.01488 (0.01485; 0.01491)	0.05100 (0.05087; 0.05113)	0.00711 (0.00708; 0.00713)
1st order autocorr.	0.15342 (0.14964; 0.15720)	0.40094 (0.39885; 0.40303)	0.39838 (0.39562; 0.40115)	0.42836 (0.42574; 0.43098)
2nd order autocorr.	-0.12215 (-0.12637; -0.11793)	-0.24278 (-0.24566; -0.23990)	0.03476 (0.03174; 0.03778)	-0.05605 (-0.05958; -0.05253)

After identifying a region of the parameter space in which the model gives rise to reasonable macroeconomic dynamics, several parameters are fine-tuned in order to replicate the moments and statistics shown in table 1 with the simulated time series. The values of many parameters turn out to be quite similar to those shown in Delli Gatti and Reissl (2020) while others differ somewhat. The full set of parameter values is shown in table 5 in appendix C.

The model is simulated 100 times with different seeds for a duration of 756 periods, equivalent to 63 years. Simulated monthly time series are transformed into annual ones and then filtered in order to construct the simulated equivalents of the empirical moments and statistics. The simulated moments are shown in table 2 as means across Monte Carlo runs along with the associated confidence intervals.

As in Delli Gatti and Reissl (2020), the model is able to closely reproduce the empirical standard deviations of GDP and consumption. Simulated investment however is significantly more volatile than its empirical counterpart. Since simulated GDP consists solely of private consumption and investment (along with constant public expenditure for healthcare) in our model and there is no role for net exports, the empirical volatility of GDP and consumption can only be reproduced if simulated investment is more volatile than empirical investment.

Similarly, the simulated employment rate is much more volatile than in the empirical data. This is due to the simplified nature of the labour market in our model which leads to employment being tied to current production much more closely than it is in reality.

Table 2: Simulated Moments & Statistics

Statistic	GDP	Consumption	Investment	Employment rate
Std. deviation	0.01564 (0.01466; 0.01662)	0.01524 (0.01432; 0.01616)	0.09254 (0.08738; 0.09769)	0.01536 (0.01441; 0.01632)
1st order autocorr.	0.59272 (0.56285; 0.62259)	0.58364 (0.55346; 0.61381)	0.47737 (0.44081; 0.51393)	0.59020 (0.56015; 0.62025)
2nd order autocorr.	0.07775 (0.03137; 0.12412)	0.05158 (0.00415; 0.09901)	0.04948 (0.00680; 0.09217)	0.07511 (0.02866; 0.12157)

The model does reasonably well at reproducing most of the first order autocorrelations but performs less well on second and higher-order autocorrelations especially regarding GDP and consumption. Similarly to Assenza et al. (2015) and Delli Gatti and Reissl (2020), we plot the autocorrelations of output, consumption, investment and the employment rate

up to lag 6 in figure 1, while figure 2 shows the cross-correlations of output, consumption, investment and the employment rate with output.

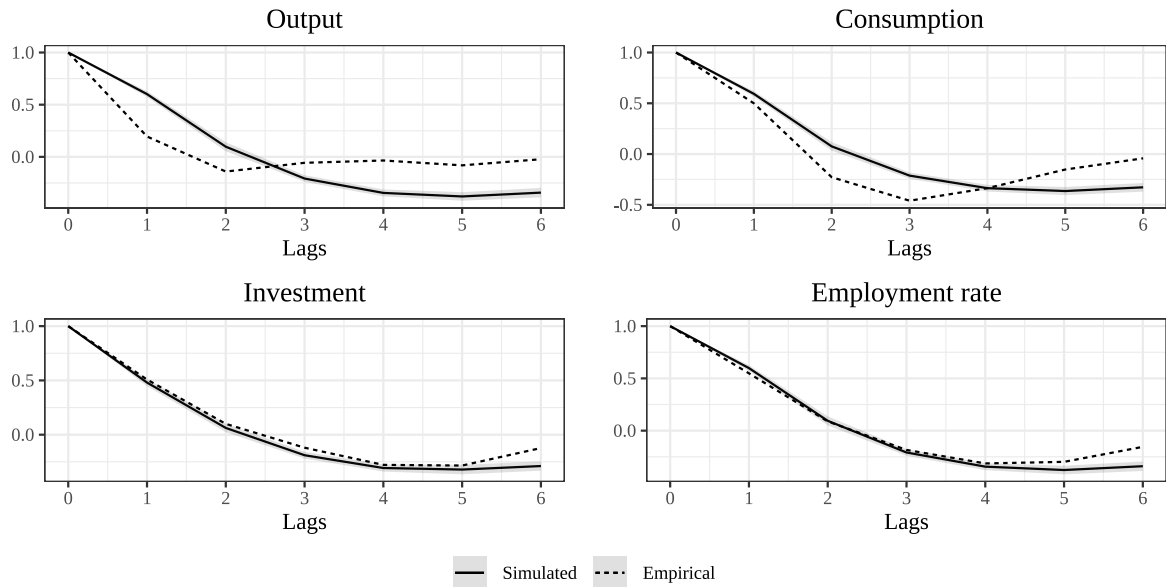


Figure 1: Empirical and simulated autocorrelations

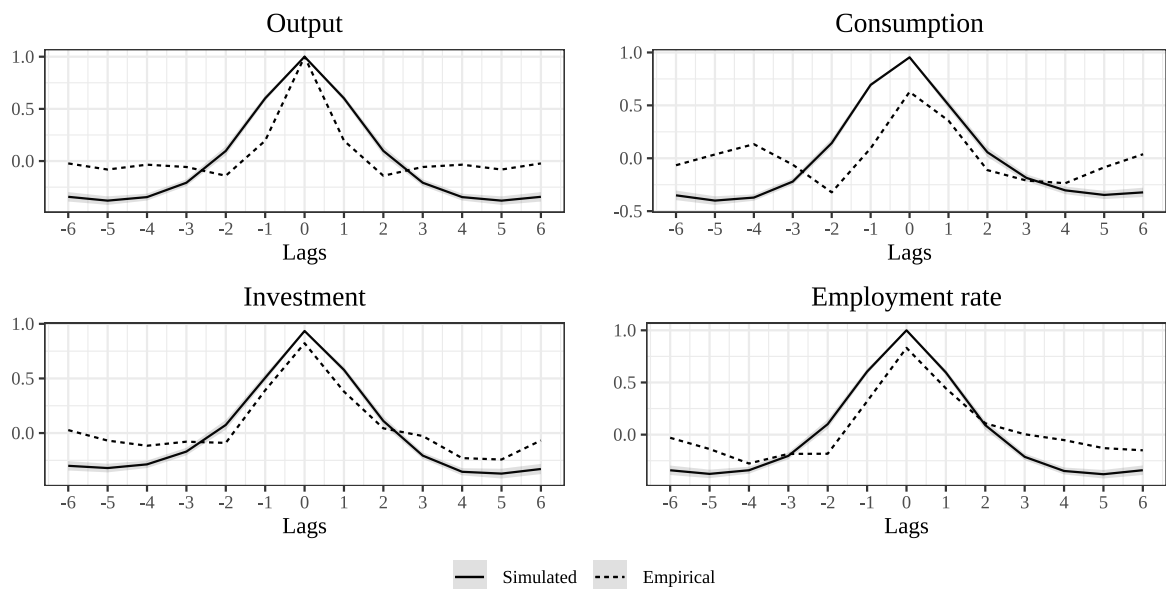


Figure 2: Empirical and simulated cross-correlations

5 Epidemic scenarios

In this section we analyse the epidemiological dynamics produced by the scenarios also analysed in Delli Gatti and Reissl (2020). The next section will be devoted to effects of the epidemic on macroeconomic variables. In what follows, we denote with t_E the *month* and τ_E the *week* of that month in which the epidemic begins. In addition, let τ_L^s and τ_L^e

denote the weeks in which the lockdown begins and ends respectively. All the scenarios described below begin by changing the status of a small number of agents from *susceptible* to *infected* in τ_E . The disease then initially spreads slowly before infections begin to pick up.

The **Uncontained Epidemic scenario (UE)** considers a case in which the infectious disease breaks out but this does not lead to any change in the behaviour of agents or to any countermeasures (Non-Pharmaceutical Interventions, NPI hereafter) taken by the government, e.g., a lockdown. Under these circumstances, the disease can spread freely throughout the population. This scenario is of course completely hypothetical, but provides a benchmark against which the effects of behavioural changes or government measures can be evaluated.

The (voluntary) **Social Distancing scenario (SD)** denotes the case in which agents spontaneously engage in “social distancing”. As in Delli Gatti and Reissl (2020), social distancing is modelled as a discrete choice by individual agents (see also Baskozos et al., 2020). We define a *distancing index* d_τ which is governed by the law of motion:

$$d_\tau = \iota d_{\tau-1} + (1 - \iota) \mathbf{N} \times \mathbf{D}_\tau, \quad (1)$$

where \mathbf{N} is a row vector containing three parameters describing the intensity of choice and \mathbf{D}_τ is a column vector containing three indicators influencing agents’ decision to distance. The first is a measure of the severity of the epidemic, given by $\mathcal{D}_{c,\tau} - \bar{\mathcal{D}}_{c,SD}$ where $\mathcal{D}_{c,\tau}$ is the number of *currently infected and detected* individuals and $\bar{\mathcal{D}}_{c,SD}$ is a fixed threshold value.⁵ The second captures social influence and is given by $\phi_d - \phi_{nd}$, that is the difference between the share of agents which are already socially distancing (ϕ_d) and those who are not (ϕ_{nd}). The third is a perceived cost of social distancing, denoted by c_{SD} which is multiplied by -1 . The probability for an agent to distance is then given by

$$\pi_\tau^d = \frac{1}{1 + \exp(-d_\tau)}, \quad (2)$$

meaning that it is increasing in the index d and asymptotically tending to 1. In words: an agent is more likely to distance (i) the higher the number of currently detected cases, (ii) the higher the fraction of agents who are already distancing, and (iii) the lower the perceived cost of social distancing. In each period we draw a random number n_h from a uniform distribution $\mathcal{U}(0, 1)$ for each agent and if $\pi_\tau^d > n_h$, the agent will engage in social distancing. Social distancing has a range of important effects on both the economic and the epidemiological dimensions of the model.

⁵In the simulations we set $\bar{\mathcal{D}}_{c,SD} = 5$.

First of all, SD affects the probability of infection conditional on an encounter between a susceptible and an infected individual. Recall that in the (uncontained) epidemic, a fraction (transmission rate) of meetings between susceptible and infected agents leads to a new infection with probability 1. Under SD, when an infectious agent i meets a susceptible agent j , we assume that the meeting between these agents generates an infection with probability

$$\pi_{i,j}^d = 1 - \beta \mathbf{1}_{\pi_\tau^d > n_i} - \beta \mathbf{1}_{\pi_\tau^d > n_j} \quad (3)$$

where $\mathbf{1}_{\pi_\tau^d > n_z}$ is an indicator function which takes value 1 if agent z is socially distancing and 0 otherwise, $z = i, j$.⁶ For each of these meetings, therefore, the probability goes down to $1 - \beta$ if one of the agents involved is distancing, and $1 - 2\beta$ if both are distancing. Second, SD reduces the overall number of connections. In the presence of social distancing, for all $t > t_E + 2$, the number of social connections and shop connections is reduced by $(1 - \phi_d)$.⁷

Third, the SD scenario is also associated with an increased effort to detect asymptomatic cases. Recall that in the (uncontained) epidemic scenario the detection probability (π^r) is exogenous. Under SD the detection probability becomes endogenous and time varying. When $t > t_E + 2$, under SD the detection probability in week τ is given by $\pi_\tau^r = \min(\overline{\pi^r}, \underline{\pi^r} + \gamma_d \dot{\mathcal{D}}_{\tau-1})$, where $\dot{\mathcal{D}}_{\tau-1}$ is the number of newly detected cases in the previous week and $\overline{\pi^r}$ is an upper bound. In the lockdown scenario (LD), the detection probability follows the same law of motion once $\tau > \tau_L^e$.

Last but not least, social distancing affects agents' demand for consumption goods. The first time any agent socially distances, their demand for L-goods receives a negative shock while their demand for B-goods receives a positive shock. The shocks are calibrated such that in percentage terms, the demand for luxury goods declines more strongly than that for basic goods increases. The shocks are given by

$$\begin{aligned} s_B &= 1 + \sigma_\tau^B \mathcal{D}_{c,\tau} \\ s_L &= 1 - \sigma_\tau^L \mathcal{D}_{c,\tau}, \end{aligned} \quad (4)$$

where

$$\begin{aligned} \sigma_\tau^B &= \max(\underline{\sigma}^B, \overline{\sigma}^B * \exp(-z * \hat{\tau})) \\ \sigma_\tau^L &= \max(\underline{\sigma}^L, \overline{\sigma}^L * \exp(-z * \hat{\tau})), \end{aligned} \quad (5)$$

and $\overline{\sigma}^B < \overline{\sigma}^L$ as well as $\underline{\sigma}^B < \underline{\sigma}^L$. In the absence of a lockdown $\hat{\tau}$ is equal to 0 (as long

⁶To assess whether agent z actually distances we draw a random number from a uniform distribution with unit support. The z -th agent engages in social distancing if $\pi_\tau^d > n_z$ where n_z is the number drawn pertaining to z .

⁷The number of connections, however, cannot go below a lower bound given by the exogenous share of connections which take place during the lockdown, described below.

as $t \leq t_E + 2$), and then increases by 1 in each week. The rationale is that the size of the consumption shock gradually declines over time as households become accustomed to the new epidemic environment.

In the **Lockdown scenario (LD)**, alongside voluntary social distancing there is also a (government mandated) lockdown. The lockdown comes into force in week τ_E^s , when the number of newly detected infections reaches the (exogenous) threshold⁸ \dot{D}_{lock} and ends in week τ_L^e . If the situation does not improve, the LD ends after a maximum duration of d_{max}^{lock} weeks. If the situation improves prior to d_{max}^{lock} – i.e., if the average of newly detected cases over the previous 2 weeks falls below a floor \dot{D}_{end} – the lockdown is lifted.⁹ We assume that the lockdown is enforced only once. In other words, if a lockdown has been imposed and subsequently lifted, there are no subsequent lockdowns even if detected infections rise beyond the threshold again in future periods. The lockdown triggers the following events in the model:

- At the beginning of the month following the institution of the lockdown, a fraction of firms producing luxury goods (L-firms) are shut down completely (and their production is halted).¹⁰ They become unable to sell any goods they may have already produced during that period. All firms which are closed in lockdown lose all of the customers who are using them as their “go-to” firm. Firms producing basic goods and capital goods (B-firms and K-firms respectively) continue producing.
- All firms immediately move into “smart working” meaning that a subset of workers of each firm work from home.
- The lockdown eliminates part of the connections at the workplace: the number of connections at every firm (in smart working) is reduced by a fixed factor (uniform across firms). The exact subset of connections which does take place is sampled anew in every period. The fraction of L-firms which are closed completely do not give rise to any workplace connections during the lockdown.
- The lockdown limits social gatherings, eliminating a part of the connections in the social network of agents. Hence the number of social connections in the LD scenario is a fraction of the corresponding number in normal times. Connections occurring in the shopping network are also reduced by the same factor. This may be interpreted as agents making fewer shopping trips than they otherwise would and, for instance, increasingly relying on deliveries of goods.¹¹ In addition, we assume

⁸In our simulations we set $\dot{D}_{lock} = 30$ newly detected cases.

⁹In our simulations $d_{max}^{lock} = 12$, i.e., the maximum duration of the lockdown is 3 months. The lockdown is lifted if the average of new detected cases over the previous 2 weeks falls below $\dot{D}_{end} = 12.5$.

¹⁰In the simulations we assume that one third of L-firms are shut down during the lockdown.

¹¹In the simulations we assume that only one fourth of social and shop connections survives during the lockdown.

that the lockdown lowers the perceived cost of social distancing, making it more likely, *ceteris paribus*, that any individual agent will engage in social distancing.

- Once the lockdown has ended, previously closed firms are re-opened but remain in smart working mode. Expected demand is re-initialised for all L-firms to account for the re-entry of the reopening firms.¹² After the lockdown, each firm moves out of smart working after a stochastic number of periods. Encounters between agents slowly adjust back to their previous level, as does the perceived cost of social distancing.

This lockdown policy is designed to mimic, in a stylised way, the policies implemented by the Italian government during spring of 2020 which, in addition to restrictions of social contacts and mandated remote working, also involved the temporary closure of a range of economic sectors deemed as “inessential”, including in manufacturing and non-customer-facing services. The vast majority of these restrictions had been lifted by June 2020. A second, ‘softer’ set of lockdown measures based on the division of the country into different zones with different levels of restrictions was implemented during autumn 2020 to combat the second wave (cf. Reissl et al., 2021; Ferraresi et al., 2021). For simplicity we do not implement this second set of lockdown measures in the model but instead assume that after the end of the (first) lockdown and from the second wave onward, only voluntary social distancing stands in the way of the epidemic disease. The LD scenario can hence be considered to mimic the actual institutional setting in Lombardy up to the start of the second wave and a counterfactual scenario from that point onward.

The effects of the lockdown on the network of connections between agents is illustrated using figures 3 and 4 which show an example of the network (encompassing all three types of connections, i.e., workplace, shops and social) during one period in normal times and one in lockdown. The reduced number of connections is immediately obvious, and is also demonstrated by an examination of two simple measures of connectivity. The network depicted in figure 3 has a density of 0.0021 and the largest eigenvalue of the corresponding adjacency matrix is equal to 66.11. In figure 4, by contrast, the density has declined to 0.00027 and the largest eigenvalue is equal to 25.05.

¹²Expected demand for each previously open L-firm is set to the minimum between the mean production of open L-firms in last period and that firm’s own production in last period. Re-opening L-firms’ expected demand is set to equal the mean production of open L-firms in the previous period. Firms’ expectations regarding capacity utilisation are adjusted in line with this.




Figure 3: Network of agents during normal times




Figure 4: Network of agents under lockdown

In figure 5 we show the simulated epidemic curves at weekly frequency. We run the model 100 times with different random seeds and compute the mean of the simulated data for each period along with 95% confidence intervals. The top left panel shows the number of cumulative infections in the three scenarios while the top right panel shows the flow of new infections. The bottom panels show the same curves for detected cases. Figure 6 instead compares the dynamics of deaths, showing both cumulative and new deaths. In all cases, the numbers have been scaled by a factor of $\frac{1}{0.003}$ in order to transform infection and death number emerging from our model with a population of 30000 into equivalents for the case of Lombardy which has a population of around 10 million.

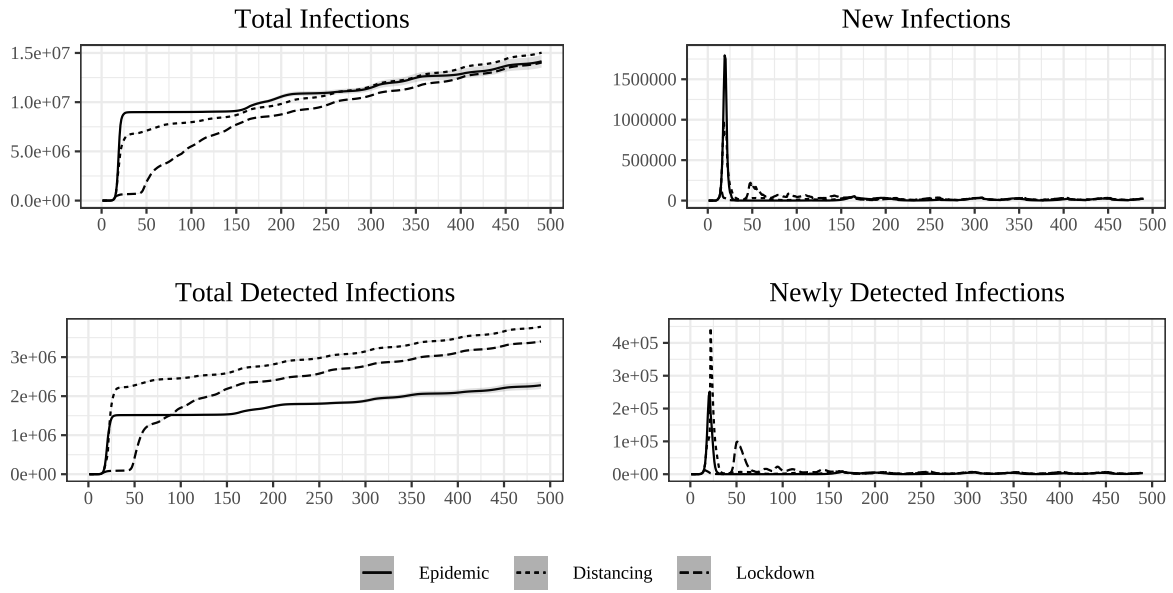


Figure 5: Comparing the epidemic scenarios (weekly)

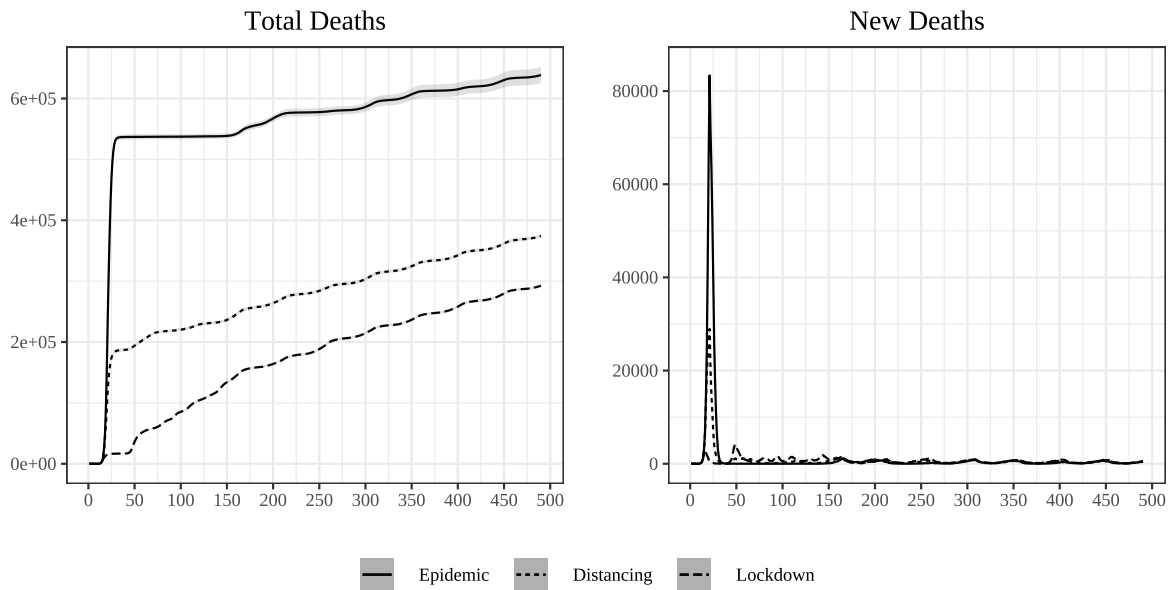


Figure 6: Comparing the epidemic scenarios (weekly)

In the UE scenario, cumulative infections grow extremely fast before reaching a plateau (let's label it the “herd immunity” number of infections) around week 25 (top left panel). At the plateau most of the (surviving) agents have been infected and have therefore left the population of *susceptibles*. In the absence of sufficient hosts, infections almost disappear for a prolonged period of time (approximately until week 150). In terms of flows, we observe a first wave of per-week new cases characterized by remarkable amplitude. Herd immunity comes at the price of a very high number of deaths as shown in Figure 6. In 3 out of 4 simulation runs herd immunity is not sufficient to eradicate the disease. In the present model, in fact, the population of susceptibles increases over time due to the

newly born and to the temporary nature of immunity: previously infected and recovered people may become susceptible again in any future period with a small probability. Once a sufficient number of agents have become susceptible again, a second wave will occur and, with the passing of time, additional but milder – in terms of duration and amplitude – waves of the disease will break out, further increasing the number of infections and deaths.¹³

In the SD scenario, the initial trajectory is similar to that of UE but the exponential trend is broken slightly earlier such that the cumulative number of infections at the plateau is lower than in the UE case. The flow of new infections per period, therefore, reaches a lower peak than in the UE scenario. The smaller number of infections in and of itself naturally also translates into a smaller number of deaths. In addition, the reduced strain on the healthcare system relative to the UE case due to the lower number of infections leads to a further reduction of deaths.

Total *detected* infections in the SD case exceed those detected in the UE case after the first wave due to the assumed positive correlation between the detection probability and the flow of new detected infections in the SD scenario. Due to the lower number of cumulative actual infections, however, in the SD scenario the second wave breaks out earlier than in UE as there is a smaller number of agents who have acquired natural immunity. In addition, these dynamics are driven by the seasonality of the base transmission rate. Once the effect of social distancing has caused the first wave to abate, the share of agents engaging in social distancing declines, enabling infections to grow again. Like the UE scenario, the SD case is subsequently characterised by a succession of milder waves as the disease becomes endemic. In contrast to the UE case, the disease does not die out in any of the 100 runs we perform of the SD scenario.

In the LD scenario, the adoption of mandated lockdown measures is able to break the first wave much earlier and at a considerably lower level of cumulative infections and deaths than in the UE and SD scenarios. In our model, the lockdown is hence highly effective at containing the first wave of the disease, hence also preventing a large number of deaths. The lagged adjustment of social and workplace connections following the lockdown as well as the remaining effects of voluntary social distancing are then able to contain new infections at a low level for some time until gradual relaxation together with the assumed seasonality of the base transmission rate lead to the emergence of a second wave. Since we assume that there is no second lockdown in this scenario, and since there is a much smaller number of agents who have acquired natural immunity, this second wave is in fact more severe than the first one. The LD scenario then continues for some time to produce slightly higher infection numbers than the SD case until the two eventually line up and produce very similar-looking dynamics

¹³The previous version of ABC (see Delli Gatti and Reissl (2020)) did not feature either newly born or reinfections so that there was only one wave.

Figure 7 compares the empirical dynamics of cumulative detected infections and per-week detected infections in Lombardy to the simulated epidemic curves from the LD scenario for the first year of the epidemic, beginning in calendar week 9 of 2020 (first week of March: this is week zero in the diagram). Overall, the model does a good job at reproducing both cumulative and newly detected infections throughout the first wave, although it appears to slightly under-estimate the number of infections taking place between the first and second waves. The model also correctly reproduces the timing of the outbreak of the second wave. Once the second wave has started, simulated infections naturally peak at a higher level than in the empirical data as we do not model the second set of lockdown measures.

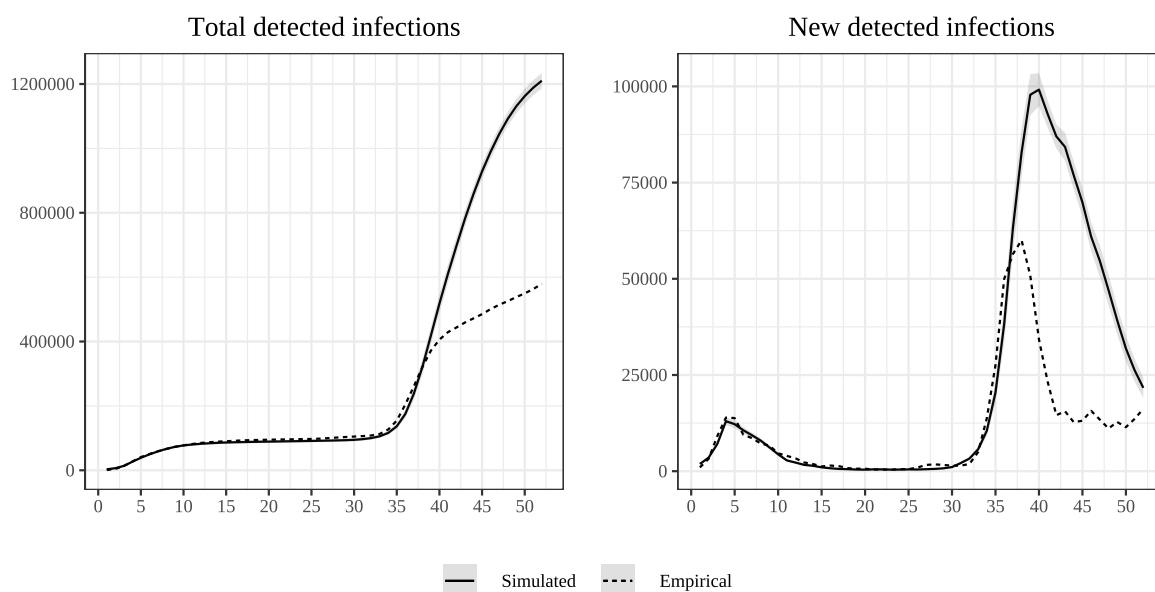


Figure 7: Comparing empirical and simulated infection data (weekly)

Figure 8 compares the simulated and empirical numbers of cumulative deaths for the same time-period as that shown in figure 7. While the model does a good job at reproducing the cumulative number of deaths at the end of the first wave, it can be seen that during the first and particularly during the second wave, simulated deaths increase prior to their empirical counterparts. This is due to the fact that in the model, deaths lag infection numbers by fewer periods than in the real world: for a given capacity of the healthcare system, a patient who develops serious symptoms in the model is as likely to die during the first week in which they are ill as in the last whereas in the real world, fatalities due to Covid-19 typically take place considerable time after the contraction of the disease.

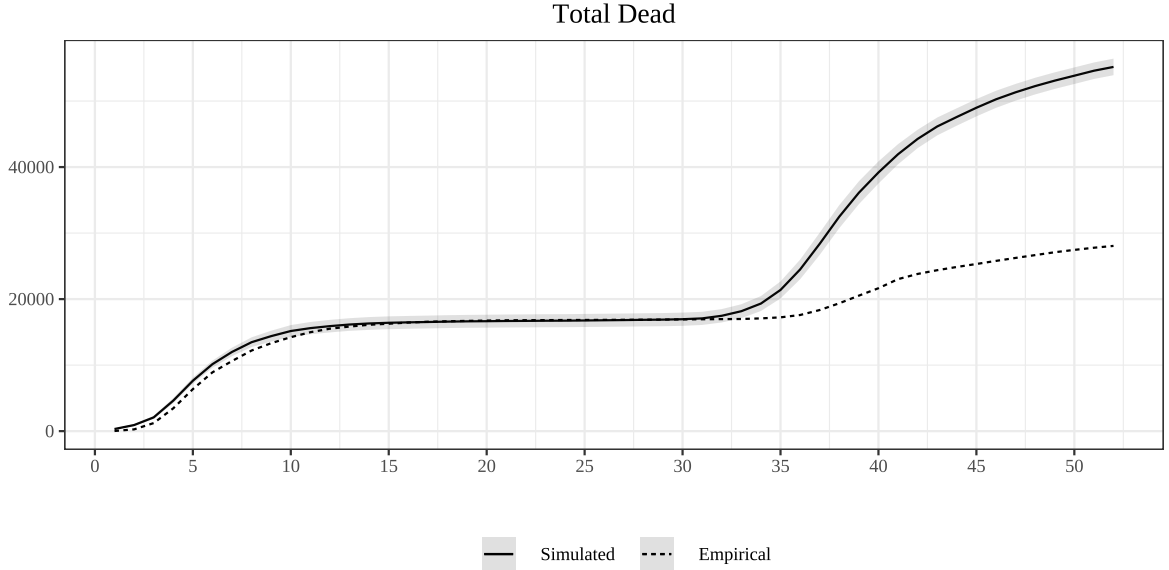


Figure 8: Comparing empirical and simulated deaths (weekly)

Table 3 compares the actual numbers of cumulative deaths and detected infections in Lombardy in calendar week 39 of 2020 to the simulated numbers 30 weeks after the outbreak of the disease in the model. Once again it can be seen that the lockdown scenario comes very close to reproducing the empirical numbers.

Table 3: Cumulative epidemic data 30 weeks after outbreak

	Simulated	<i>Lower</i>	<i>Upper</i>	Empirical
Total Detected	94593	90034	99153	104752
Total Deaths	17473	16459	18488	16922

6 The macroeconomic effects of the epidemic

In this section, we briefly examine the macroeconomic effects of the epidemic under the same scenarios presented in the previous section. As we said, we run the model 100 times with different random seeds under the scenarios of uncontained epidemic (UE), social distancing (SD) and lockdown (LD). For a given run r , for each macroeconomic variable *in levels* – say x – we compute the *percent deviation* of the variable in scenario $i = UE, SD, LD$ from the baseline scenario of Normal times (NT): $\hat{x}_t^{r,i} = \frac{x_t^{r,i} - x_t^{r,NT}}{x_t^{r,NT}}$. Then we compute the mean of this percent deviations across 100 runs: $\hat{x}_t^i = \sum_{r=1}^{100} \hat{x}_t^{r,i}$. For variables that are already expressed *in percent terms* (government debt to GDP and the bankruptcy rate) – say y – we compute the *absolute deviation* $\Delta y_t^{r,i} = y_t^{r,i} - y_t^{r,NT}$ and the mean of the absolute deviations $\Delta y_t^i = \sum_{r=1}^{100} \Delta y_t^{r,i}$. The time series of these

means – i.e., \hat{x}_t^i and Δy_t^i , $t = 0, 1, \dots, 120$ – are plotted in figures 9 and 10 (along with 95% confidence intervals). Month zero is the first month of the epidemic.

We begin by examining the UE scenario. Recall that in this setting, it is assumed that agents do not alter their behaviour in response to the epidemic and there are also no measures taken by the government to contain the epidemic. As such, this experiment should not be viewed as a plausible empirical scenario; instead it serves as a purely hypothetical case to be compared to the more plausible settings of voluntary social distancing, coupled with or in the absence of a government-mandated lockdown. In the UE scenario, GDP, consumption and investment all decline during the initial phase of the epidemic and remain below their pre-epidemic levels for the entire period plotted in the figures. This decline is chiefly due to the reduction in the size of the population. As illustrated in the previous section, the UE scenario leads to a high number of deaths. While dead agents are replaced in the model, this process is very gradual such that the population remains well below the baseline for an extended time in the UE case. In addition to reducing potential aggregate output, the consumption expenditure coming from the people who die from the disease is “removed” from aggregate demand as long as they are not replaced. The lower level of overall economic activity also leads to a lower price level relative to the baseline. The relative price of luxury goods, on the other hand, is hardly affected in the UE scenario. The default rate initially increases slightly due to the decline in economic activity but then returns toward the baseline. While transfers to households for income support increase in the short term due to the high number of agents who fall ill, transfers for pensions decline permanently as most of the agents who die from the disease are inactive. With unchanged tax rates, tax revenue also declines due to some economically active agents dying from the disease. The decline of pension outlays, however, more than offsets the short-term increase of benefit payments and the reduction of tax revenues. Hence, government debt declines both in absolute terms and, eventually, as a share of GDP.

The SD scenario gives rise to a sharp contraction in GDP that is much more pronounced than the one produced by the UE setting. This decline is chiefly driven by a collapse in aggregate consumption, due to the assumed effect of social distancing on consumption demand. Since social distancing on its own is not able to contain infections as effectively as when coupled with a lockdown, the high resulting infection numbers make the negative shock to consumption larger than in the LD scenario. Firms respond by reducing their investment, and the bankruptcy rate increases in the short run due to firms’ sales declining sharply. The price level at first declines, before increasing and slightly overshooting during the recovery of consumption, the latter being driven by the smaller numbers of infections in subsequent waves as well as the assumed decline in the size of the consumption shock over time. As the disease becomes endemic, GDP and consumption remain depressed for an extended period since the effect of infections on consumption demand continues

to exert an influence. This in turn also appears to lead to a permanently lower price level. Since the shock to consumption demand is assumed to also feature a substitution effect from luxury toward basic goods, the relative price of luxury goods initially declines strongly and then increases for some time as the demand for luxury goods rebounds. The economic downturn leads to a sizeable increase in government debt both in absolute terms and as a share of GDP.

Recall that in the LD scenario, one third of all L-firms are shut down and cease to produce any output during the lockdown. This large supply shock is immediately reflected in the aggregate data, leading to a sharp decline in GDP. In addition, firms which are forced to close are unable to sell the output they have already produced, leading to a spike in the bankruptcy rate. There is also a decrease in consumption, partly driven by social distancing and partly by reduced output of luxury goods, as well as in investment. While social distancing leads to a decline in the demand for luxury goods, the supply of these goods decreases even more strongly due to mandated closures during the lockdown. This leads to an increase in the relative price of L-goods which in turn also drives an increase in the price level. Recall that in the LD scenario (and in contrast to the SD case), due to the successful suppression of the first wave of the disease the second wave of the epidemic ends up being more severe than the first. Due to the high number of infections during the second wave, there is hence a second large shock to consumption after a brief recovery period, which drives a renewed decline in GDP and investment. GDP recovers slowly following the second wave and then overshoots for some time. This overshooting of GDP (along with consumption and investment), which is also present to a smaller degree in the SD scenario, is due to the adaptive rules of thumb underlying agent behaviour in the model. Investment in particular overshoots the baseline quite strongly due to firms rebuilding capacity lost during the previous recession. In addition, the rule of thumb which firms use to make their investment decision leads them to over-react to both positive and negative changes in capacity utilisation. Eventually, GDP returns to a level slightly below the baseline due to the disease having become endemic. The LD scenario also leads to a sizeable increase in government debt, exceeding that observed for the SD scenario.

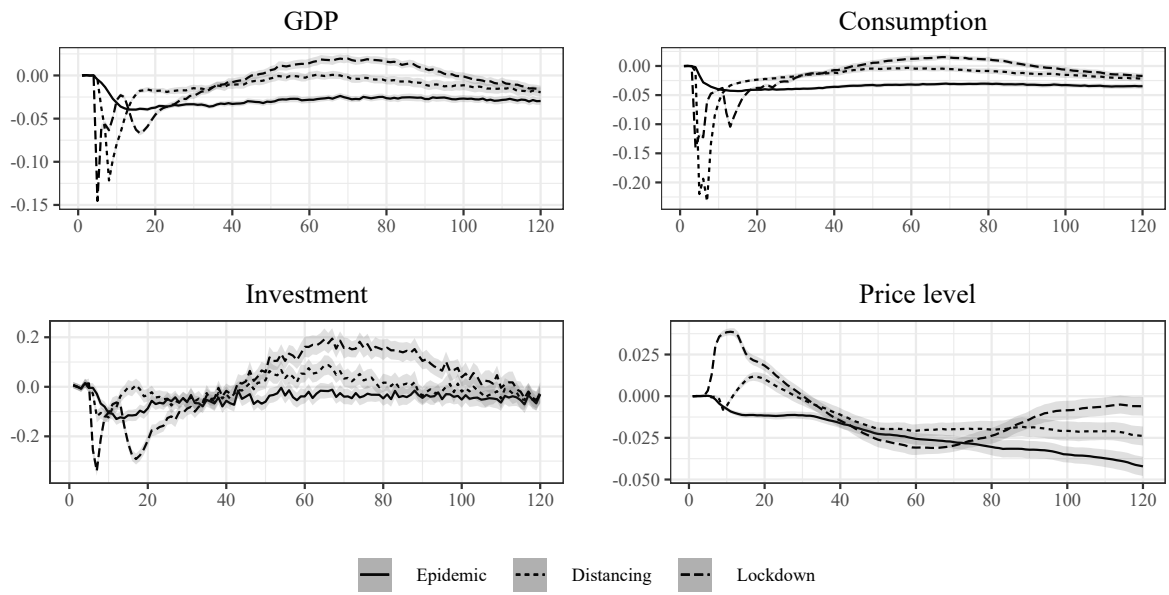


Figure 9: Economic impact of the disease under different scenarios

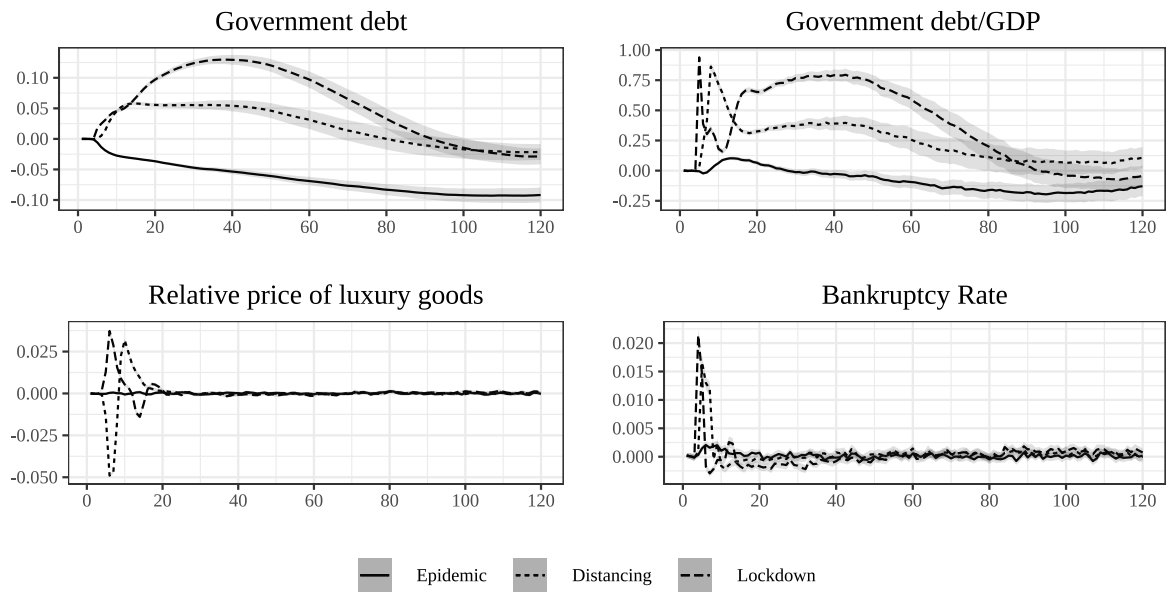


Figure 10: Economic impact of the disease under different scenarios

The first panel of figure 9 shows the percentage deviation of GDP in absolute terms from its value in the baseline. Since in the various epidemic scenarios, the population of the model is not constant, it is useful to also consider the dynamics of GDP *per capita* relative to the baseline, shown in figure 11.

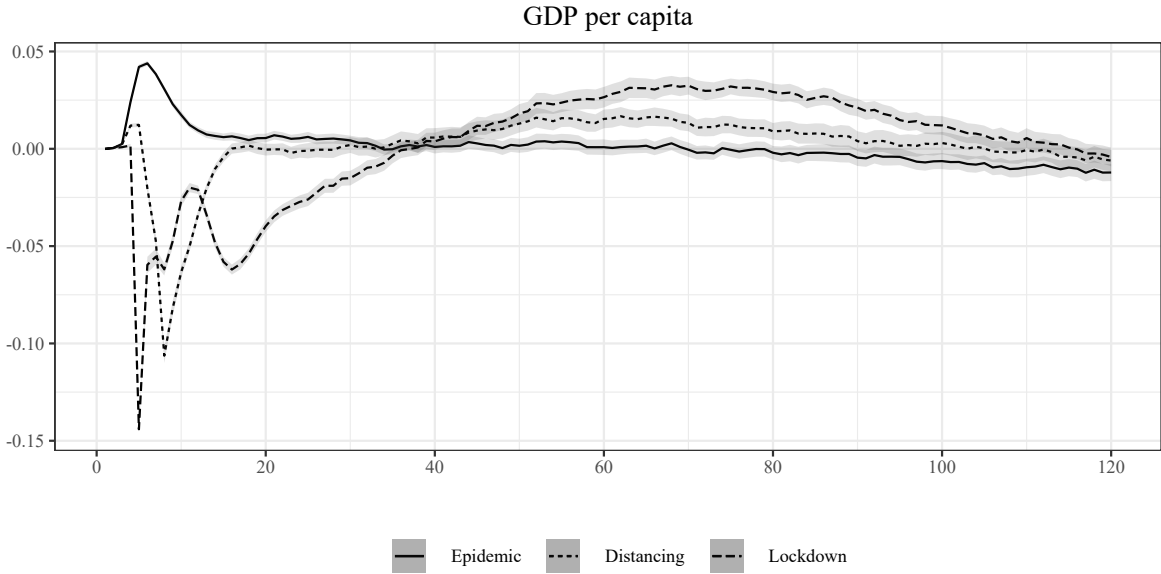


Figure 11: Economic impact of the disease under different scenarios

It can be seen that in the UE scenario, while GDP decreases in absolute terms, GDP *per capita* in fact exhibits an increase, i.e. the decline in population outweighs the decline in economic activity, particularly in the early stage of the epidemic. This result can be explained by the fact that the epidemic primarily leads to the deaths of economically inactive agents who do not contribute to the productive capacity of the model economy and whose consumption demand financed by pension income is lower than that of an economically active and employed agent. This result does not, however, make the UE scenario desirable even from a purely economic viewpoint as the increase in per capita GDP does not necessarily imply that any individual economically active agent is better off; instead it is purely driven by the fact that a share of the population which on average contributes less to GDP is removed. Beyond economic considerations, recall that the UE scenario implies an extremely high number of deaths and moreover, as stressed above, should be considered as a purely hypothetical case which we discuss only for comparative purposes.

In the SD setting, the dynamics of GDP per capita initially look very similar to those of GDP in absolute terms, but GDP per capita returns all the way back to its baseline value after the initial decline, whereas GDP in absolute terms remains below the baseline for a more extended period. As above, this effect can be explained by the distribution of fatalities across the population.

In the LD scenario, the dynamics of GDP per capita closely resemble those of GDP in absolute terms, suggesting that population dynamics do not play a large role in driving the macroeconomic aggregates in this setting.

In all three scenarios, it can be seen that toward the end of the time horizon depicted in figure 11, GDP per capita appears to tend to a level slightly below its baseline value, with

this tendency being strongest in the UE case and weakest in the LD case. The reason for this is that the epidemic, by overwhelmingly leading to the deaths of old agents who are eventually replaced by young ones, leads to a demographic change. Pension payments are permanently lower, while tax rates are unchanged, such that the government budget balance tends to increase. This in turn represents a slight but noticeable drag on aggregate demand which leads to a lower GDP. Due to the high number of deaths compared to the SD and especially the LD case, this effect is strongest in the UE scenario.

7 Vaccination

Having analyzed the effects of the epidemic disease under different scenarios, the current and following sections address the economic and epidemiological impacts of vaccination, both in the presence and absence of variants of the epidemic disease. In what follows, we take the LD scenario as the new benchmark against which the effects of vaccination and variants will be assessed.

7.1 Vaccine efficacy

As in reality, immunity to the epidemic disease in our model can take two forms: natural immunity is acquired by means of infection and recovery; vaccine-induced immunity is generated by vaccination. Vaccination may have three distinct effects: it may reduce the susceptibility of vaccinated individuals to infection, it may reduce the exposure of vaccinated individuals to serious disease (and hence mortality), and it may reduce the transmissibility of the infection by a vaccinated individual who has become infected. Vaccine Efficacy may hence be measured with respect to:

1. susceptibility (VE_1)
2. lethality (VE_2)
3. transmissibility (VE_3)

With regard to vaccines against Covid-19, there is ample evidence that vaccine-induced immunity of type 1 is not complete: the estimates vary from around 60 to around 90% ($60\% < VE_1 < 90\%$) depending on the vaccine (lower for traditional vaccines, higher for mRNA vaccines). Hence, a vaccinated person may still be infected, but with a lower probability. To capture this feature, we assume that when an infectious agent i meets a susceptible agent who has received the vaccine v , the probability π^t that i will transmit the disease to v is reduced by $0 < VE_1 < 1$. The probability of transmission in case of vaccination, therefore, is described by a modified equation (3), namely:

$$\pi_{i,v}^t = (1 - VE_1 \mathbf{1}_{\tau,v}) (1 - \beta \mathbf{1}_{\pi_i^d > n_i} - \beta \mathbf{1}_{\pi_v^d > n_v}) \quad (6)$$

where $\mathbf{1}_{\tau,v}$ is an indicator function which takes value 1 if v is vaccinated and 0 otherwise. Clinical evidence suggests that the existing vaccines' ability to prevent serious disease (VE_2) is sizable and somewhat higher than VE_1 ; estimates range from 90 to 95%. In the model, the subset of agents who will develop serious symptoms if infected is defined at the beginning of the simulation.¹⁴ Agents in this subset will develop serious symptoms with a probability equal to 1 in the absence of vaccinations. We assume that if an individual that is part of this subset is vaccinated, they will develop serious symptoms only with probability $1 - VE_2$.

On the basis of the evidence collected so far, it is unclear whether vaccines are effective also in reducing the probability that a vaccinated individual who has nevertheless got Covid-19 infects other people. Estimates of VE_3 vary considerably across studies. We set ourselves in the worst case scenario assuming that a vaccinated infected individual is as infectious as an unvaccinated one.¹⁵ Therefore we set $VE_1 = 0.7$; $VE_2 = 0.9$ and $VE_3 = 0$.

We assume that a single dose of the vaccine provides an immunity, the duration of which (in weeks) is given by a random draw from a normal distribution with mean equal to 52. In other words, we assume that vaccine-induced immunity lasts on average for one year. Thereafter the protection disappears ($VE_1 = VE_2 = VE_3 = 0$) and the individual goes back to the status of susceptible. We rule out for simplicity the presence of people who refuse to get vaccinated: all individuals eligible for vaccination (i.e., susceptibles, recovered and *undetected* infected) accept the vaccine when it is offered to them.

¹⁴The shares of agents (in each age class) who will develop serious symptoms are denoted with π_z^h ; $z = y, m, o$. The numerical values of these parameters are in table 6. Agents who are newly born during the simulation fall in the category of young people.

¹⁵As to VE_1 , data submitted by Janssen to the FDA, from a randomized control trial for their single-dose vaccine showed a vaccine efficacy of type 1 of 67% when considering cases occurring at least 14 days after the single-dose vaccination. In December 2020, a study on the effects of Moderna found that vaccination implied a lower risk of symptomatic Covid-19 of 94%. For the Pfizer BioNTech vaccine the most remarkable example of its efficacy are the results obtained in Israel through the mass vaccination campaign that began in December 2020. The Clalit Institute for Research, in a study that took place between December 2020 and February 2021, found that two doses of Pfizer vaccine reduced symptomatic cases by 94%. On March 22 2021 a report containing the results of the Phase III trial of the AstraZeneca vaccine in the US showed a vaccine efficacy at preventing symptomatic COVID-19 of 76%, that occurred 15 days or more after receiving two doses with an interval of four weeks between the two. In addition, results were comparable across age groups, with an increased vaccine efficacy of 85% in adults 65 years and older. As to VE_2 , the efficacy of Moderna in terms of lower risks of hospitalisation due to severe Covid-19 symptoms was estimated to be of 89%. Pfizer reduced the risk of hospitalisation by 87% and severe Covid-19 symptoms by 92%. With regard to preventing severe or critical disease and hospitalisation, the Astra Zeneca vaccine demonstrated 100% efficacy. These data are collected with reference to the original virus. However they seem to apply also to the variants emerged so far. As to VE_3 at the time of writing (August 2021), evidence from the US points to setting the transmissibility of the virus by vaccinated individuals to the same level of that of non-vaccinated individuals. This seems to suggest no efficacy of type 3, at least with respect to the delta variant.

7.2 Vaccination strategies: three scenarios

There are two key features which characterize a vaccination policy:

- the *coverage rate*, i.e., the fraction of the population which can be vaccinated in one unit of time;
- the *priority rule*, i.e., the procedure defining which target groups (if any) should be prioritized for vaccination, e.g. old or young.

Both dimensions can affect the efficacy of the vaccination campaign, the evolution of the epidemic and the resulting effects on the macroeconomy.

We define the coverage rate as the share of the population which can be vaccinated *in every week*. This may be influenced by vaccine availability as well as the technical and organizational capacity of the vaccination authorities, both of which tend to improve over time. In the simulations, in order to replicate the actual time-line in Italy, we assume that the vaccine is introduced 11 months after the outbreak of the epidemic (corresponding to December 2020) – i.e., well after the end of the first lockdown. We also assume that the coverage begins at a low level (0.01) and then increases by 0.001 in every period until reaching an upper bound (0.05). The pool of agents eligible for the vaccine consists of susceptible, recovered and undetected infected agents. The order in which agents are selected from the pool in a given period is defined by the priority rule. The latter may be designed in a way to place more weight to certain conditions than others, say health-demographic versus economic factors, depending on the particular strategy adopted by the policy maker.

We compare three alternative vaccination strategies based on different priority rules: Randomized Vaccination (RV); Priority by Age (PA); Priority to Workers (PW). In the RV scenario, agents are randomly sampled from the pool of eligible individuals. In the PA scenario, the probability of being drawn from the pool increases (exponentially) with age, thus giving priority to the old. Finally, in the PW scenario, the probability of being drawn is highest for economically active workers, that is young and middle-aged, and lowest for economically inactive, i.e., the old.¹⁶ As the vaccination campaign progresses, however, the strategies become essentially equivalent because the number of available doses per week will eventually exceed the number of unvaccinated agents.

From an epidemiological point of view, we would expect vaccination to reduce both the number of infections which occur every week, as well as the mortality of the disease.

¹⁶In the PA scenario, the sample weights are given by $exp(1)$, $exp(2)$ and $exp(3)$ for young, middle aged and old respectively, such that the probability of being drawn for vaccination increases exponentially with age. Since all the old are inactive by assumption – meaning that the labour force consists only of young and middle aged agents – in PW the weights are given by $exp(1)$ for the old and $exp(3)$ for the young and middle-aged.

From an economic point of view, we would therefore expect vaccination to reduce the duration and amplitude of the output loss due to the epidemic, and consequently also the resulting public sector deficit and debt. Moreover, our framework allows to gain insights about the effects of alternative vaccination strategies both at the macro level, in terms of aggregate health and economic outcomes, and at the meso level, by comparing the number of infected and fatalities between age groups.

7.3 Economic and epidemic effects of vaccination

Figure 12 compares the simulated epidemic curves under the different vaccination strategies to the baseline, that, in this section, is the LD scenario without vaccination. As in the previous experiments, we run the model 100 times with different random seeds and compute the mean and 95% confidence intervals for each time period. In the figure we consider a time window consisting of 100 weeks (which correspond to more than 2 years in our simulations)¹⁷ starting from the beginning of the vaccination campaign. The left panel shows the cumulated number of (detected) infected individuals computed *from the beginning of the epidemic* which would have occurred with and without the vaccine. The right panel shows the first difference of this number, i.e., the per week flow of new infections. Since vaccination started well after the end of the (first) lockdown in Italy (June 2020), the time series shown in the figure start when the lockdown and the first wave have already ended. The waves shown in the right panel are therefore the (latter part of the) second and subsequent ones.

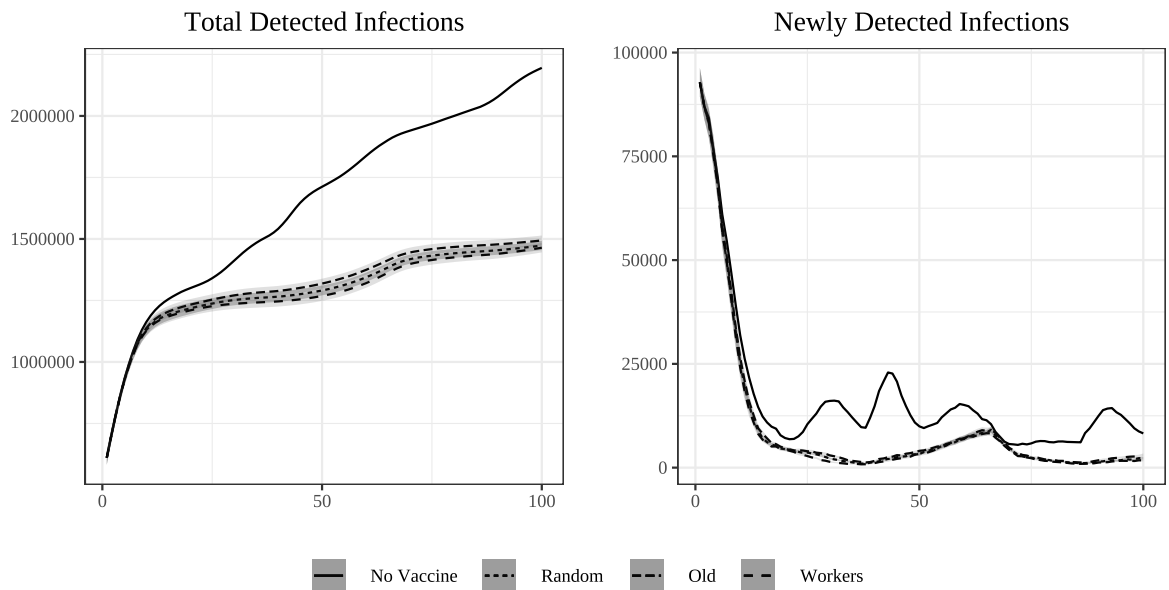


Figure 12: Impact of vaccination by strategies (weekly)

¹⁷We assume that there are 4 weeks per month so that a year consists of 48 weeks in the simulation.

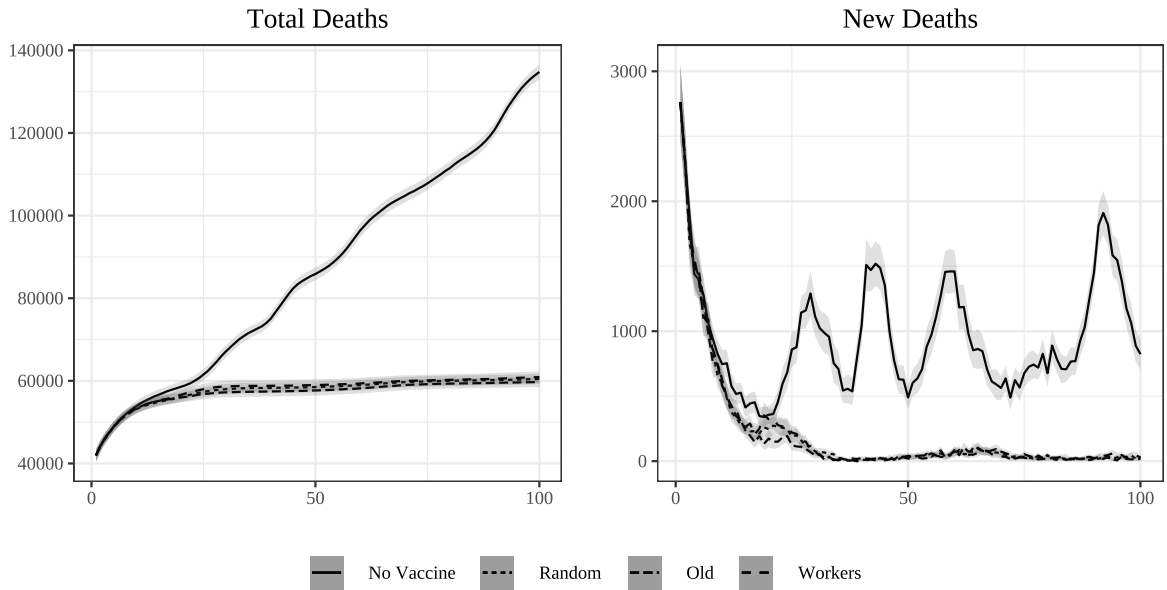


Figure 13: Impact of vaccination by strategies (weekly)

From the left panel, it can be seen that the vaccine significantly contributes to reducing contagion. Given the axis scale, the differences among prioritization strategies are not appreciable to the naked eye. Regardless of the prioritisation strategy adopted, however, one year after the start of the vaccination campaign the cumulative number of infected individuals in the presence of a vaccine is about 400 thousand units (approximately 23%) lower than in the absence of it. Moreover, the gap between cumulative detected infections with and without vaccine is widening over time. As shown in the right panel, the introduction of the vaccine slightly accelerates the decline of the second wave.¹⁸ In the absence of a vaccine, there will also be multiple subsequent waves, endogenously determined by reinfections, albeit much less pronounced than the second one. With the vaccine, in the time window considered we observe only one additional wave which starts approximately one year after the beginning of vaccination and peaks six months later at a much lower level than the second wave. This is because the vaccine-induced immunity is not complete, thus leaving vaccinated agents partially exposed to the risk of contagion. *The vaccine acts as a mitigator* leading to a reduction both of the frequency and the amplitude of subsequent waves of contagion.

As expected, the decline in the cumulated number of infections due to vaccination, coupled with the vaccine's ability to prevent serious disease, triggers a significant reduction in fatalities, as shown in figure 13. The left panel shows the cumulated number of deceased individuals since the beginning of the epidemic which would have occurred with and without the vaccine in the same time window as in figure 12 while the right panel shows the first difference of this number, i.e., the weekly flow of new deaths. The vaccine

¹⁸Note that, as already stated, we do not have a second lockdown in place in our simulations. Hence the number of infections in the simulations are much higher than in reality.

significantly contributes to saving lives. From the left panel we infer that, regardless of the prioritisation strategy adopted, one year after the start of the vaccination campaign the number of deceased individuals stabilizes while it keeps growing (with oscillations) in the absence of a vaccine. One year after the start of the vaccination campaign the cumulative number of deceased individuals in the presence of a vaccine is about 27 thousand units (approximately 30%) lower than in the absence of it. As shown in the right panel, after the second wave, in the presence of the vaccine the curve of new deaths remains substantially flat at zero while there are multiple endogenous waves of deceased in the absence of a vaccine. The reason of this disconnect between the dynamics of infections and deaths lies in the fact that the vaccine efficacy at preventing serious symptoms is higher than that at protecting from contagion ($VE_1 < VE_2$).

From the figures above it is difficult to discern a clear difference in the epidemic dynamics across vaccination strategies. Table 4 takes a snapshot of the epidemiological situation after the first year of the vaccination campaign (from week zero to week 48). It shows the total number of deaths and detected infections (averaged across simulations) which occurred *during the first year of the vaccination campaign*¹⁹ under different priority rules, for the whole population and by different age groups. First and foremost, the table confirms that, independently of the prioritization strategy, vaccination leads to a sizable improvement in the epidemiological situation. Comparing the level of each epidemiological group (in the lockdown scenario) in the absence of a vaccine with the average level across vaccination strategies in the presence of a vaccine, we observe a large reduction in the number of detected cases, and an even more dramatic one in the death toll.

From the same table we infer that prioritization by age group with priority given to the old allows to save more lives – in particular among the old for obvious reasons – at the cost of a somewhat higher number of infections. On the other hand, prioritization by economic activity with priority given to active workers leads to lower infections but higher fatalities. The reason for this is quite straightforward. Contrary to old and inactive agents who meet only with social and marketplace connections, employed workers also interact with colleagues. Therefore, a vaccination strategy aimed at prioritizing workers, by protecting individuals with greater connectivity, reduces the overall level of contagion, but at the same time leaves the elderly, i.e., the subjects with the highest likelihood of developing serious symptoms, more exposed to the risk of dying.

¹⁹In the left panels of figures 12 and 13 we show the cumulative numbers of detected infections and fatalities since the beginning of the epidemic. The numbers in table 4 are not comparable with those used to produce the figures, since the former do not include detected infected or deceased from the beginning of the epidemic to the beginning of the vaccination campaign.

Table 4: Number of deaths and infections in the 1st year, with and without vaccination (average across simulations; extremes of the 95% confidence interval in parentheses)

	No Vaccine	Random	Old	Workers
Total Deaths	43927 (42932, 44922)	16553 (15746, 17360)	15737 (14896, 16577)	17040 (16174, 17906)
Total Detected	1091697 (1055116, 1128277)	676730 (644038, 709422)	704047 (672028, 736065)	654537 (622000, 687074)
Dead (young)	477 (399, 554)	283 (223, 344)	243 (184, 303)	203 (152, 255)
Dead (middle)	12147 (11688, 12605)	5820 (5443, 6197)	6257 (5855, 6658)	5970 (5570, 6370)
Dead (old)	30813 (30060, 31567)	10413 (9848, 10979)	9197 (8640, 9754)	10847 (10189, 11504)

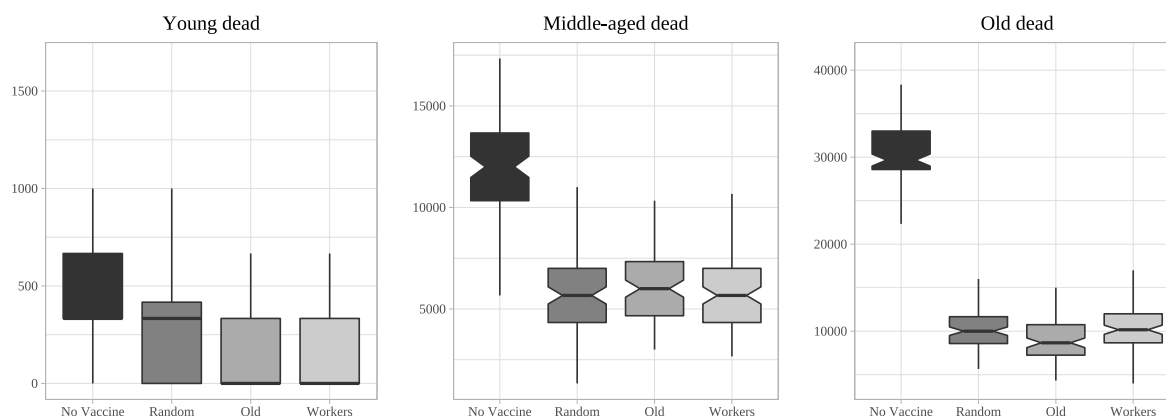


Figure 14: Differences in fatalities across age groups under alternative vaccination strategies

The box plots in figure 14 provide a clearer idea of the differences in fatalities across age groups under alternative vaccination strategies. Once again, the differences between different priority rules are more pronounced among the elderly, who benefit most from an age-oriented priority rule. The within-group variation between strategies is very limited among the young and middle-aged since the absolute number of fatalities among these age groups is relatively small across all vaccination scenarios.

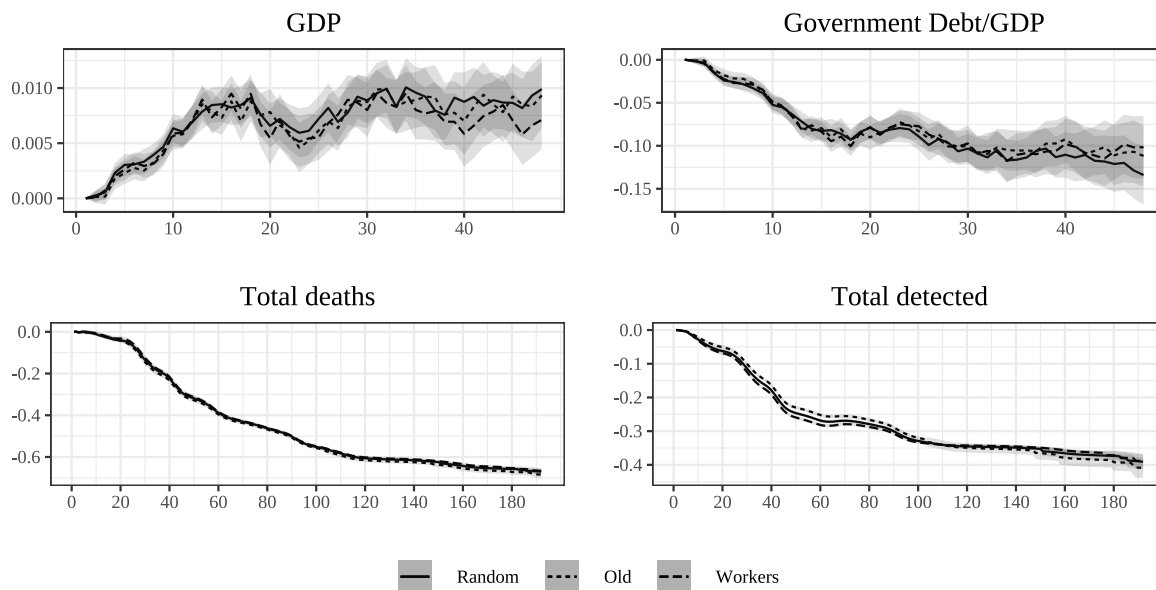


Figure 15: Impact of vaccination by strategies (monthly (top) and weekly (bottom))

Figure 15 shows the deviation of GDP, government debt as a share of GDP (both at monthly frequency), cumulative deaths and detected infections (both at weekly frequency) under different vaccination strategies from the lockdown scenario with no vaccine, starting from the first period of the vaccination campaign. The introduction of the vaccine leads to a small but significant and permanent effect on GDP which after one year of vaccination campaign amounts to nearly 1%. This is essentially due to the positive impact on consumption and aggregate demand resulting from the reduced number of deaths and infections brought about by the diffusion of the vaccine-induced immunity. In particular, the lower number of infections reduces the probability to engage in social distancing as well as the size of the (negative) consumption shock associated with it. In addition, the lower number of deaths results in higher aggregate consumption because the consumption of the old who survive due to the vaccine is not temporarily “removed” until they are replaced by young agents. This boost in GDP also leads to a reduction in government debt as a share of GDP relative to the lockdown scenario.

8 Variants

Finding a vaccine is certainly the most effective way to combat the epidemic, but it may not be sufficient to eradicate it. Viruses, in fact, evolve rapidly to ensure their own survival. To do so, some pathogens mutate the shape of their proteins, to avoid being targeted by antibodies that would prevent them from infecting cells (Zimmer, 2021). This has happened to the SARS-CoV-2 virus and has immediately become a matter of concern for scientists and governments.

A succession of variants of SARS-CoV-2 has emerged around the world over the course of

the Covid-19 pandemic.²⁰ At the time of writing (August 2021), according to the Italian National Institute of Health, almost all infections in Italy are with the Delta variant. Variants may enhance the infectiousness of the virus and/or its lethality. One feature that variants so far have in common is enhanced transmissibility, while evidence on changes in lethality to date is mixed (also depending on the particular variant). Importantly, the variants appear to have an impact on vaccine effectiveness, both in preventing infection and serious disease.

8.1 Variant types, timing, spreading

To capture these diverse aspects, we model two types of variant. Both variants are more contagious than the original virus, as reflected in a higher transmission rate ρ_c (+50%) and a lower effectiveness of endogenous social distancing β (-75%); the fatality rate, π^m , instead, is left unchanged. The variants differ with respect to their degree of resistance to the original vaccine. While Variant 1 is able to reduce only the vaccine’s ability to prevent infection – i.e., VE_1 – Variant 2 reduces both VE_1 and VE_2 , i.e., the vaccine’s efficacy in preventing serious disease which may eventually lead to death. The vaccine in our model may however be adapted over time to be fully effective against the variant, in which case agents need to be re-vaccinated with the new vaccine. Accordingly, for each variant we study two alternative scenarios depending on whether or not a new vaccine becomes available.

The variant is injected into the model in period τ_V , a few weeks before the start of the vaccination campaign – the timing in the model therefore approximately mimics the time at which the Alpha variant was detected in the United Kingdom – through a small set of “super-spreader” agents \mathcal{I}_{v,τ_V} , i.e., households with a high number of connections to susceptible agents. This group constitutes the agents who are initially infected with the variant and spread it throughout their network of interactions. At first the variant can coexist with the original virus, implying that susceptible agents may be infected with either one of the two. Given the higher transmission rate, however, the variant becomes dominant fairly quickly and endogenously replaces the original virus. In the scenario featuring an improved vaccine, the new vaccine becomes available 24 weeks after the emergence of the variant, meaning that it takes half a year to adapt the old vaccine and make it available to the public. Once available, the new vaccine is first administered to any agents who are completely unvaccinated and subsequently to those who had previously received the old vaccine.

²⁰At the time of writing, there are as many officially detected variants as there greek letters between α and λ . The major identified variants “of concern” are: B.1.1.7 (Alpha) and B.1.351 (Beta) initially detected in September 2020 in the United Kingdom and in South Africa respectively; P.1 (Gamma) and B.1.617 (Delta) initially detected in December 2020 in Brasil and in India respectively.

The diffusion of Variant 1 is illustrated²¹ in figure 16.

In the figure we consider a time window consisting of 200 weeks (which correspond to more than 4 years in our simulations) starting from the week in which the epidemic begins (week zero). The baseline in this section is the PA scenario, i.e., the LD scenario augmented with vaccination in the case of priority given to the old. The solid line in the left panel shows the flow of per-week new infections due to both the original virus (beginning in week zero) and the variant (emerging approximately one year later) in the absence of an improved vaccine. The dotted line represents the flow of individuals infected with the variant. The right panel shows the same variables in the presence of an improved vaccine (which is introduced approximately half a year after the variant emerges). The first wave in the figure therefore is due exclusively to the original virus. The variant emerges just after the peak of the second wave and after 12 weeks accounts for nearly 30% of new cases (as shown by the dashed line). It takes less than half a year for the variant to become dominant and replace the original virus.

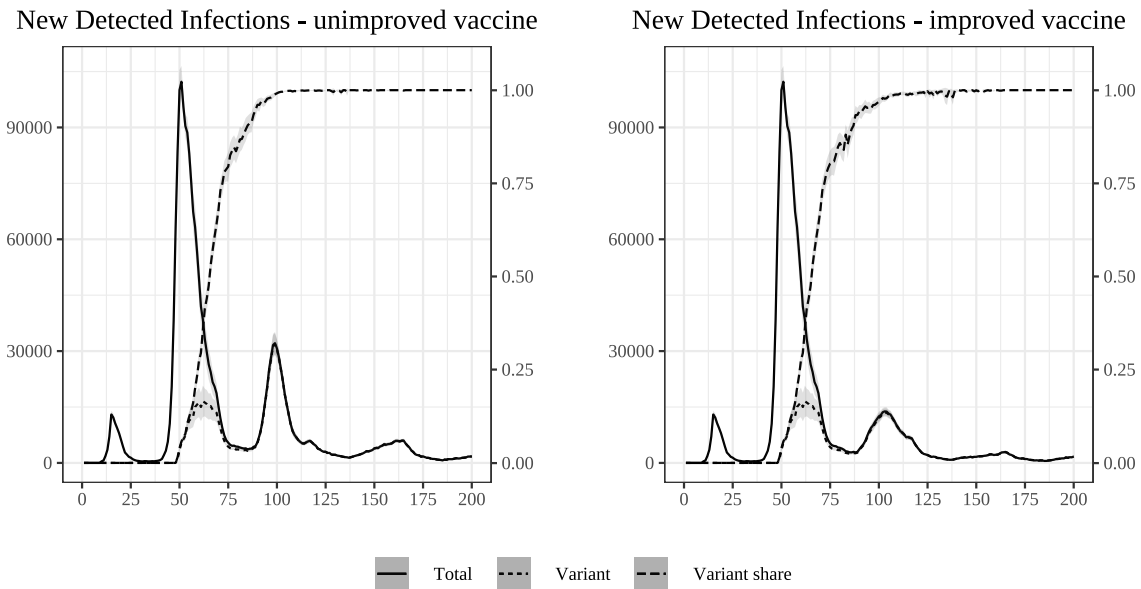


Figure 16: Spreading of Variant 1 with and without improved vaccine

The *variant acts as an accelerator* leading to an increase of the frequency and amplitude of the waves succeeding the second one. Therefore, it partially offsets the effects of the vaccine. If the vaccine is improved and adapted to the variant, as shown by the right panel, the amplitude of the subsequent waves is mitigated again.

8.2 Economic and epidemic effects of variants

In what follows, we analyze the effects of the emergence of Variant 1 and Variant 2, with and without an improved vaccine, on epidemic and economic dynamics. Figure 17

²¹Since it has the same transmission rate, the pace of diffusion of Variant 2 is broadly similar.

compares the simulated epidemic curves under different combinations of vaccinations and variants to the baseline, that, in this section, is the vaccination scenario with the age-oriented priority rule where only the original virus circulates. In the figure we consider a time window consisting of 100 weeks starting from the week in which the original vaccination campaign begins, i.e. just after the peak of the second wave. The left panel shows the cumulative number of (detected) infected individuals (computed since the beginning of the epidemic) which occur with and without the variant and with or without an improved vaccine in this time window. The right panel shows the weekly flow of newly detected infections.

From figure 17 it is clear that the emergence of Variant 1 leads to a significant increase in the number of infections due to a slower decline of the second wave and a stronger upsurge in new cases during the third wave. The negative impact of virus mutation is partly contained by the introduction of the new vaccine, which takes six months to be produced and another six months to fully unfold its effects. When infections rebound during the third wave, the new vaccine significantly reduces the peak of weekly detected infections, though infection numbers are still much higher than in the absence of a variant.




Figure 17: Impact of Variant 1 with and without improved vaccine

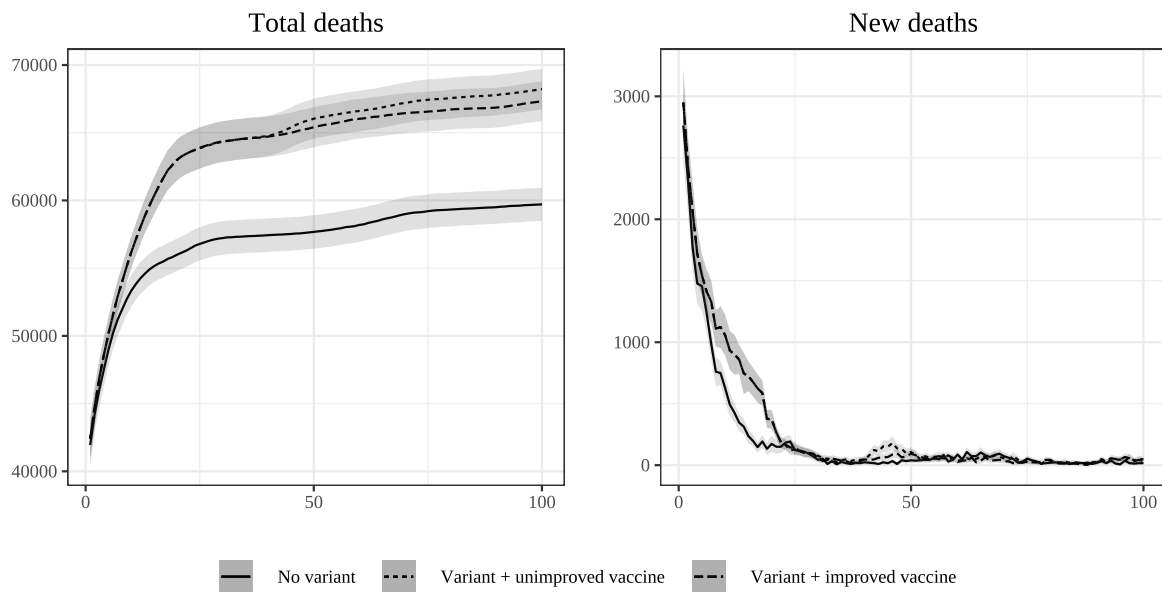


Figure 18: Impact of Variant 1 with and without improved vaccine

Despite the variant not being inherently more lethal than the original virus, the increased number of infections also leads to a higher number of deaths, as shown by Figure 18. The left panel shows the cumulative number of deceased individuals (computed since the beginning of the epidemic) which occur with and without the variant and with or without an improved vaccine in the same time window as in Figure 17. The right panel shows weekly new deaths.

Within one year from the emergence of the variant the number of deceased increases by around 8 thousand units compared to the scenario without a variant. Interestingly, the improved vaccine does not appear to have a significant impact on fatalities. This is because Variant 1 does not affect the vaccine's ability to prevent serious symptoms (VE_2); it only reduces the vaccine's efficacy at protecting from infections (VE_1). Therefore, the increase in the number of deaths that we observe in figure 18 is entirely due to the higher level of contagion occurring once the variant circulates. As the share of population who receives the vaccine expands, the immunity induced by the (old) vaccine is strong enough to reduce the level of deaths, also in the presence of Variant 1 without an improved vaccine.

The situation is different in the presence of Variant 2, which undermines both dimensions of vaccine efficacy. As expected, figure 19 shows that the contagion dynamics under Variant 2 are similar to those under Variant 1. The reason is that, by assumption, the two variants have the same transmission rate, i.e., 50% higher than the original virus.

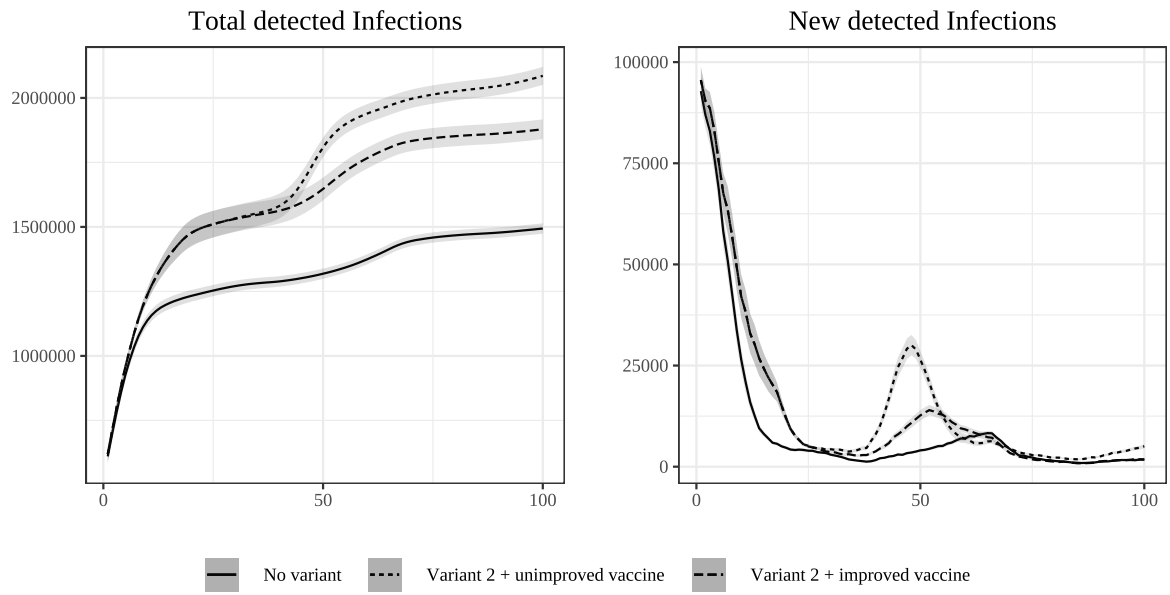


Figure 19: Impact of Variant 2 with and without improved vaccine

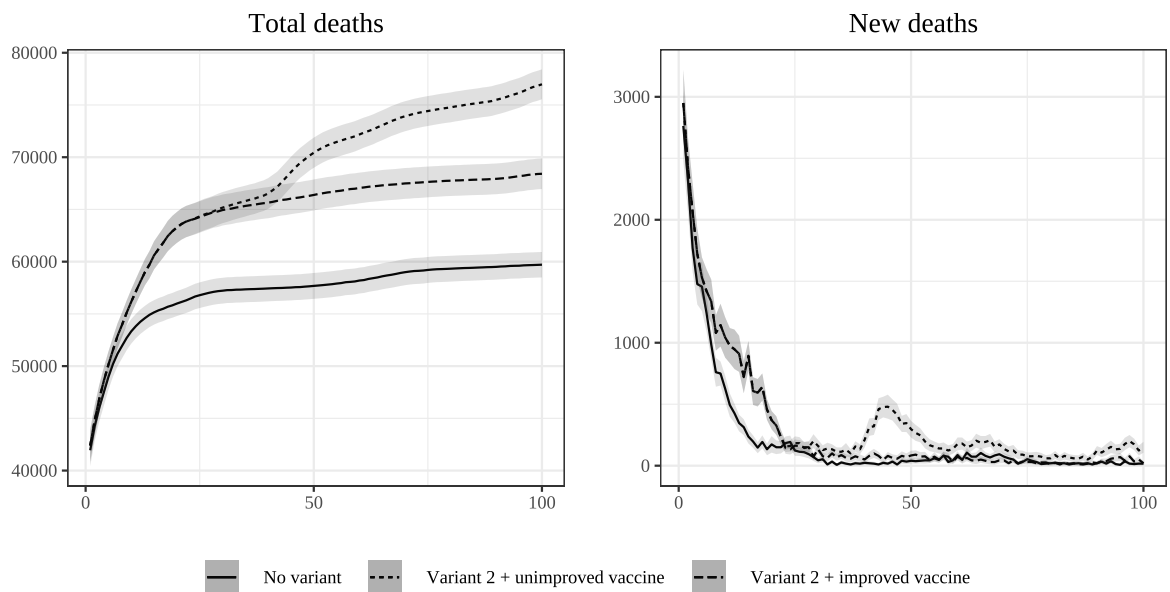


Figure 20: Impact of Variant 2 with and without improved vaccine

However, the outbreak of Variant 2, by reducing vaccine's ability to prevent serious symptoms, leads to a considerable increase in the death toll. In this case, the third wave of contagion is accompanied by a resurgence in new deaths, which can be avoided only if the healthcare sector is able to swiftly produce and distribute a new and more effective vaccine.

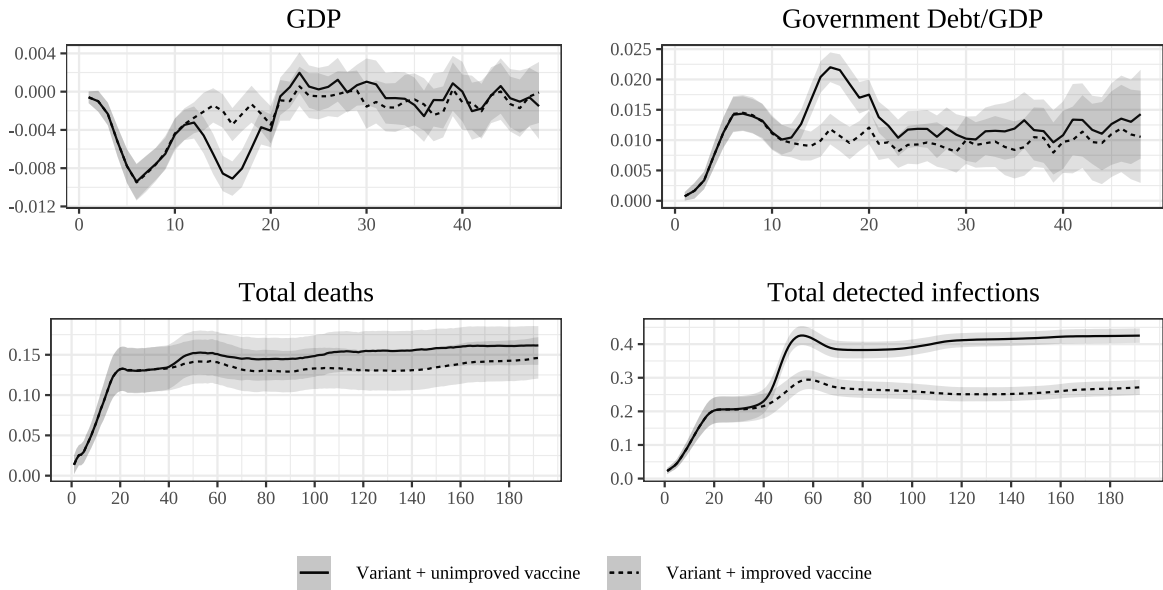


Figure 21: Impact of Variant 1 with and without improved vaccine

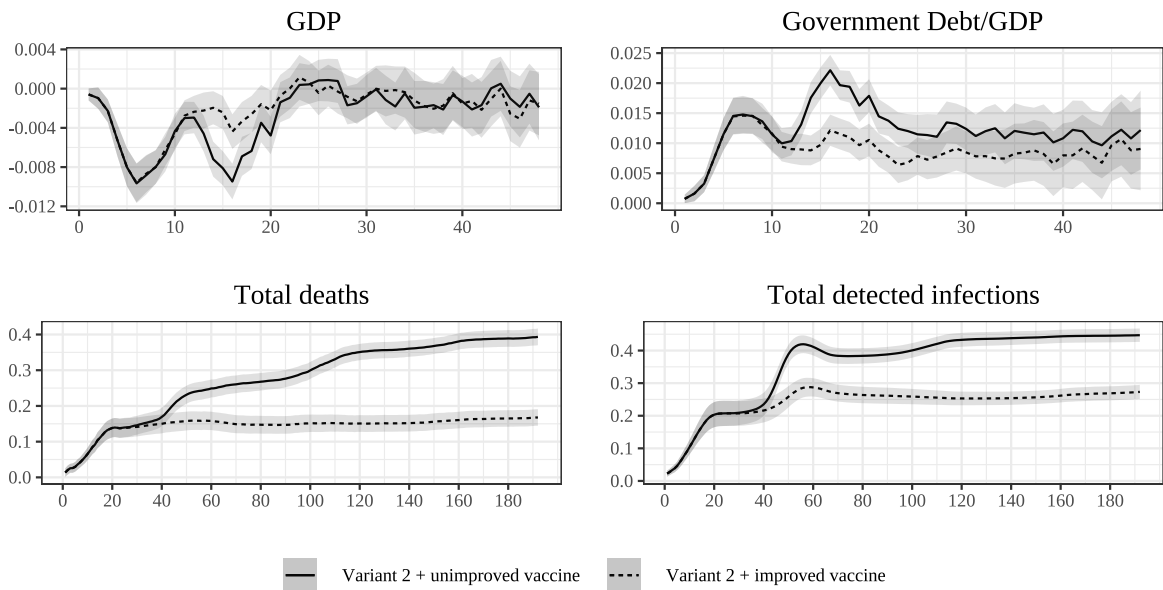


Figure 22: Impact of Variant 2 with and without improved vaccine

The macroeconomic consequences of the variants are illustrated in figure 21 and 22. The plots report the deviation of each variable from the baseline which in this case is the scenario with an age-prioritized vaccination policy is in place but without a variant.²² Under both variant types, the economy experiences an additional decline in output of nearly 0.8% during the second wave which is also reflected in the dynamics of government debt as a share of GDP, thus largely neutralizing the positive economic impact of vaccination shown in figure 15. The reduced number of infections due to the adoption

²²The deviations from the baseline are computed as described at the beginning of section 6.

of a new vaccine is able to abate the negative effects of the third wave, which otherwise would imply a further economic downturn.

9 Conclusion

In this paper we studied the epidemiological and macroeconomic effects of vaccines and variants using simulations of a macroeconomic-epidemiological agent-based model.

As expected, simulations show that other things being equal, vaccination significantly slows down the pace of the epidemic and saves lives. Over the medium run both the number of infections and of deaths drop substantially in the presence of a vaccine.

Simulations also show that vaccination is a significant *mitigating factor* of the cyclical dynamics of infections and deaths. In the absence of a vaccine, the epidemic will follow a pronounced cyclical trajectory with subsequent waves of infections and deaths. In the presence of a vaccine, both the amplitude and frequency of additional waves are decreased. We experimented with different priority rules in vaccination. Contrary to our expectations, different prioritization strategy do not translate into sizable differences in epidemiological outcomes. In our simulations, the size of the mitigation effect of vaccination on contagion is almost identical across different prioritization strategies. Priority given to the old allows to save more lives at the cost of more infections but the magnitude of this effect is relatively small.

We also experimented with two types of variants of the model virus, labelled 1 and 2. We assumed that both types are more contagious than the original virus and reduce the efficacy of the vaccine in preventing infection but variant 2 also reduces the efficacy of the vaccine in preventing serious symptoms and death.

The outbreak of variant 1 in the presence of a vaccine targeted to the original virus (a scenario that corresponds to the real world situation at the time of the outbreak of the Delta variant of Covid 19) leads to a rapid replacement of the original virus in the host population, driving a new wave of infections. Even though, by assumption, variant 1 is not more lethal than the original virus, the increasing flow of infections also leads to a significant increase in deaths. Variant 2 generates an even higher death toll because by assumption this variant is capable of also weakening the efficacy of the vaccine in preventing serious symptoms and death, sparking a third wave of deaths in the absence of an improved vaccine.

A variant hence acts as an *accelerating factor*, counteracting the mitigating effects of the vaccine until the latter can be adapted to the variant. In reality, vaccines and variants de facto co-exist and alternate in time. In fact, the probability that a variant emerges is increasing with the incidence of the disease and the availability of a vaccine against the original virus. These intertwined endogenous dynamics of vaccines and variants represent a fruitful topic for future research.

References

- ASSENZA, T., D. DELLI GATTI, AND J. GRAZZINI (2015): “Emergent dynamics of a macroeconomic agent based model with capital and credit,” *Journal of Economic Dynamics & Control*, 50, 5–28, <https://doi.org/10.1016/j.jedc.2014.07.001>.
- BABUS, A., S. DAS, AND S. LEE (2021): “The optimal allocation of Covid 19 vaccines,” *CEPR Discussion Paper*, DP15329.
- BASKOZOS, G., G. GALANIS, AND C. D. GUILMI (2020): “Social distancing and contagion in a discrete choice model of COVID-19,” *Centre for Applied Macroeconomic Analysis - Australian National University*, 35.
- BASURTO, A., H. DAWID, P. HARTING, J. HEPP, AND D. KOHLWEYER (2021): “How to design virus containment policies? A joint analysis of economic and epidemic dynamics under the COVID-19 pandemic,” *Universitat Bielefeld Working Papers in Economics and Management*, 6.
- BERNANKE, B., M. GERTLER, AND G. S. (1996): “The Financial Accelerator and the Flight to Quality,” *The Review of Economics and Statistics*, 78, 1–15, <https://doi.org/10.2307/2109844>.
- CAIANI, A., A. GODIN, E. CAVERZASI, M. GALLEGATI, S. KINSELLA, AND J. STIGLITZ (2016): “Agent based-stock flow consistent macroeconomics: Towards a benchmark model,” *Journal of Economic Dynamics and Control*, 69, 375–408.
- DELLI GATTI, D., M. GALLEGATI, B. GREENWALD, A. RUSSO, AND J. STIGLITZ (2010): “The financial accelerator in an evolving credit network,” *Journal of Economic Dynamics and Control*, 34, 1627–1650.
- DELLI GATTI, D. AND S. REISSL (2020): “ABC: An Agent Based Exploration of the Macroeconomic Effects of Covid-19,” *CESifo Working Papers*, 8763.
- FERRARESI, T., L. GHEZZI, F. VANNI, M. GUERINI, F. LAMPERTI, G. FAGIOLO, A. CAIANI, M. NAPOLETANO, A. ROVENTINI, AND S. REISSL (2021): “On the economic and health impact of the Covid-19 shock on Italian regions: A value-chain approach,” *LEM Working Paper*, 2021/10.
- FORSLID, R. AND M. HERZING (2021): “Whom to Vaccinate First - Some Important Trade-Offs,” *CEPR Discussion Paper*, DP15800, <https://ssrn.com/abstract=3784040>.
- GOLLIER, C. (2021): “The welfare cost of vaccine misallocation, delays and nationalism,” *CEPR Discussion Paper*, DP15980.
- KEOGH-BROWN, M., R. SMITH, AND J. E. A. EDMUNDS (2010): “The macroeconomic impact of pandemic influenza: estimates from models of the United Kingdom, France, Belgium and The Netherlands.” *European Journal of Health Economics*, 11, 543554, <https://doi.org/10.1007/s10198-009-0210-1>.

- LINDSKOG, M. AND I. STRID (2020): “Efficient vaccination against COVID-19 entails major macroeconomic benefits,” *Riksbank Economic Commentaries*, <https://www.riksbank.se/globalassets/media/rapporter/ekonomiska-kommentarer/engelska/2020/efficient-vaccination-against-covid-19-entails-major-macroeconomic-benefits.pdf>.
- MATRAJT, L., J. EATON, T. LEUNG, AND E. BROWN (2021): “Vaccine optimization for COVID-19: Who to vaccinate first?” *Science Advances*.
- MELLACHER, P. (2020): “COVID-Town: An Integrated Economic-Epidemiological Agent-Based Model,” *GSC Discussion Paper Series*, 23.
- REISSL, S., A. CAIANI, F. LAMPERTI, M. GUERINI, F. VANNI, G. FAGIOLO, T. FERRARESI, L. GHEZZI, M. NAPOLETANO, AND A. ROVENTINI (2021): “Assessing the Economic Impact of Lockdowns in Italy: A Computational Input-Output Approach,” *LEM Working Paper*, 2021/03.
- SAAD-ROY, C. M. AND C. E. E. A. WAGNER (2020): “Immune life history, vaccination, and the dynamics of SARS-CoV-2 over the next 5 years,” *Science*, 370, 811–818, <https://doi.org/10.1126/science.abd7343>.
- SMITH, R., M. KEOGH-BROWN, T. BARNETT, AND T. J. (2009): “The economy-wide impact of pandemic influenza on the UK: a computable general equilibrium modelling experiment,” *British Medical Journal*, <https://doi.org/10.1136/bmj.b4571>.
- ZIMMER, K. (2021): “A guide to emerging SARS-CoV-2 variants,” *The Scientist*, <https://www.the-scientist.com/news-opinion/a-guide-to-emerging-sars-cov-2-variants-68387>.

Appendix A The macroeconomic sub-model

A.1 Households

The economy is populated by N_H households, of which N_W are workers and N_F are firm owners. Households will be indexed with $h = 1, 2, \dots, N_W, N_W + 1, \dots, N_H$. Households indexed with $h \in (1, N_W)$ are workers; households indexed with $h \in (N_W + 1, N_H)$ are firm owners. Of course, the cardinality of the set of firm owners is $N_F = N_H - N_W$. As there is one owner household per firm, it coincides with the cardinality of the set of firms.

A.1.1 Workers

Workers can be economically active (employed or unemployed) or inactive (sick or retired). Each active worker supplies 1 unit of labour inelastically. If employed, they receive the uniform nominal wage w_t and pay a fraction t_w (the tax rate on wages) of this wage to the Government.

If unemployed, the worker searches for a job visiting a subset z_e of firms chosen at random. Once an unemployed worker finds a firm with an unfilled vacancy they stop searching and the match occurs. Unemployed workers who have not succeeded in finding a job receive an unemployment subsidy from the Government equal to a fraction of the wage: $s_u w_t$. A sick worker receives sick-pay $s_s w_t$. Each retired worker receives a pension $s_p w_t$. The parameters s_u, s_s, s_p are the *replacement rates* in the case of unemployment subsidy, sick-pay and pension.

A.1.2 Firm owners

The household indexed with $h = N_W + f$ is the owner of the f -th firm, $f = 1, 2, \dots, N_F$. The income of this household consists of dividends, which, in turn, are equal to a fraction ω (the pay-out ratio) of the after-tax profit $(1 - t_\Pi)\Pi_{f,t-1}$ where t_Π is the tax rate on profit and $\Pi_{f,t-1}$ is profit generated in the previous period. The firm pays out dividends only if $\Pi_{f,t-1} > 0$. If a firm faces a loss, its net worth will go down correspondingly and the firm will not distribute dividends. Moreover, the firm owners are assumed to jointly own the representative bank and consequently each firm owner receives an equal share of the dividends distributed by the bank: $\omega(1 - t_\Pi)\Pi_{b,t-1}$.

A.1.3 Households as consumers

Households receive income and interest payments $r_d D_{h,t-1}$ where r_d is the interest rate on deposits and D_h are deposits. Hence the disposable income of household $h \in (1, N_W)$

is:

$$Y_{h,t} = \begin{cases} (1 - t_w)w_t + r_d D_{h,t-1} & \text{if } h \text{ is employed,} \\ s_u w_t + r_d D_{h,t-1} & \text{if } h \text{ is unemployed,} \\ s_s w_t + r_d D_{h,t-1} & \text{if } h \text{ is sick (but not retired),} \\ s_p w_t + r_d D_{h,t-1} & \text{if } h \text{ is retired} \end{cases} \quad (7)$$

while the disposable income of household $h \in (N_W + 1, N_H)$ is

$$Y_{h,t} = (1 - t_\Pi)\omega \left(\Pi_{h,t-1} + \frac{1}{N_F} \Pi_{b,t-1} \right) + r_d D_{h,t-1} \quad (8)$$

A household's consumption decision proceeds in four steps. First, the household constructs a proxy of permanent income $\bar{Y}_{h,t}$ using an adaptive algorithm: $\bar{Y}_{h,t} = \xi_Y \bar{Y}_{h,t-1} + (1 - \xi_Y)Y_{h,t}$ where $\xi_Y \in (0, 1)$ is a memory parameter. $\bar{Y}_{h,t}$ is hence a weighted average of past disposable incomes with exponentially decaying weights.

Second, the household determines the desired *consumption budget*:

$$C_{h,t} = \bar{Y}_{h,t} + c_W D_{h,t-1} \quad (9)$$

where $c_W \in (0, 1)$ is the propensity to consume out of financial wealth, which, in this setting, coincides with deposits.

Third, the consumer allocates $c^b \in (0, 1)$ of $C_{h,t}$ to the consumption of basic goods (B-goods hereafter). Therefore $(1 - c^b)$ of $C_{h,t}$ will be devoted to purchasing luxury goods (L-goods).²³ We assume that c^b is a decreasing function of the (average) price of B-goods relative to L-goods.²⁴

Fourth, the consumer goes to the market to purchase consumption goods (C-goods) which can be either B-goods or L-goods. Consider first the market for B-goods. B-firms are indexed with $f \in (1, N_F^b)$ where N_F^b is the cardinality of the set of B-firms.

In period t , the consumer visits two firms in this set: a ‘go-to’ supplier and a potential new partner, the latter being randomly drawn from the population of B-firms (excluding the ‘go-to supplier’). The consumer then compares the prices. If the price charged by the go-to firm (say P_0) is lower than or equal to that of the potential new partner (P_1),

²³This is the *desired* allocation of the consumption budget to B-goods and L-goods. If the household's consumption budget turns out to be larger than available liquidity (deposits inherited from the past plus current income) the desired allocation will be unfeasible. In this case we assume that the consumer will first spend up to $c^b C_{h,t}$ on B-goods and then allocate any remaining liquidity to the consumption of luxury goods.

²⁴In symbols: $c^b = \frac{N_F^b P_{t-1}^l}{N_F^c P_{t-1}^b}$ where $N_F^c = N_F^b + N_F^l$ is the cardinality of the set of C-firms (the union of B-firms and L-firms) while P^b (resp: P^l) is an aggregator of the individual B-prices (L-prices). If the relative price is 1, i.e., if on average B-firms charge the same price as L-firms, the fraction of the consumption budget allocated to B-firms is $\frac{N_F^b}{N_F^c}$, i.e., it is equal to the fraction of B-firms in the population of C-firms.

they will stick to the former and will shop at the latter only if the consumption budget is not completely exhausted with the first purchase. Otherwise, they will switch to the new partner (and reverse the order of purchase) with a probability π^c which is increasing in the price set by the go-to partner relative to that of the potential new partner: $R_{0,1}^b = P_0/P_1$. In symbols:

$$\pi^c = \begin{cases} 1 - \exp \gamma_p (1 - R_{0,1}^b) & \text{if } R_{0,1}^b > 1, \\ 0 & \text{if } R_{0,1}^b \leq 1 \end{cases} \quad (10)$$

where $\gamma_p > 0$. If the consumer actually switches to the new partner, the latter becomes their new go-to partner in period $t+1$.

This *partner selection* mechanism implies an implicit *negative elasticity* of the demand for the good produced by the f -th firm with respect to the price it charges relative to that of its competitors. Consider firm f , $f \in (1, N_F^b)$. The higher is P_f relative to P_t^b , the higher the probability that the customers of the f -th B-firm will switch to a new shopping partner, reducing the demand for the f -th B-good accordingly.²⁵ If a firm goes bankrupt, all the households who have this firm as their go-to supplier will randomly choose a different go-to B-supplier.

If, at the end of their visits to B-firms, the household has not spent the consumption budget allocated to B-goods, they will save involuntarily. This market protocol does not guarantee equilibrium. Queues of unsatisfied consumers (involuntary savers) at some firms may coexist with involuntary inventories of unsold goods at some other firms.

The market protocol for L-goods follows the same lines as that of B-goods. L-firms are indexed with $f \in (N_F^b + 1, N_F^c)$ where $N_F^c = N_F^b + N_F^l$ is the cardinality of the set of C-firms (the union of B-firms and L-firms). The consumer has one “go-to” L-supplier (who sets the price P_0) and one potential partner (P_1). They will stick to the former and shop at the latter only if the budget allocated to L-goods is not completely exhausted with the first purchase in the case in which $R_{0,1}^l = \frac{P_0}{P_1} \leq 1$. They will switch to the new partner (and reverse the order of purchase) with probability $\pi^c = 1 - \exp \gamma_p (1 - R_{0,1}^l)$ if $R_{0,1}^l > 1$. If the budget allocated to L-goods has not been entirely spent, the household will add the residual to their savings.

Total saving is equal to the sum of voluntary or desired saving (i.e., the difference between disposable income and the budget allocated to consumption) and involuntary saving. This is tantamount to saying that actual saving is equal to the difference between current income and actual consumption of B-goods and L-goods. Savings are used to accumulate financial wealth in the form of deposits.

²⁵Delli Gatti et al. (2010) and Caiani et al. (2016) adopt a similar algorithm for partner selection.

A.2 Firms

There are N_F firms, of which N_F^b produce B-goods, N_F^l produce L-goods and N_F^k produce capital goods (K-goods). Firms will be indexed with $f = 1, 2, \dots, N_F^b, N_F^b + 1, \dots, N_F^c, N_F^c + 1, \dots, N_F$ where $N_F^c = N_F^b + N_F^l$. In words: firms indexed with $f \in (1, N_F^b)$ produce B-goods; firms indexed with $f \in (N_F^b + 1, N_F^c)$ produce L-goods; firms indexed with $f \in (N_F^c + 1, N_F)$ produce K-goods.

A.2.1 C-Firms

B-firms and L-firms are consumption goods producers (C-firms for short) and follow the same behavioural rules. In this section we describe the behaviour of a generic C-firm.

The firm has *market power* and sets the individual price and quantity under uncertainty. It knows from experience that if it charges higher prices it will receive smaller demand but it does not know the actual demand schedule (i.e., how much the consumers would buy at any given price). Hence the firm is unable to maximize profits since the marginal revenue is unknown. We assume that the firm charges a price as close as possible to the average price and sets a quantity as close as possible to (expected) demand. In this way the firm minimizes overproduction (in case of excess supply) or the queue of unsatisfied customers (in case of excess demand).

The f -th firm, $f \in (1 : N_F^c)$, must choose in t the price and desired output for $t+1$, i.e., the pair $(P_{f,t+1}, Y_{f,t+1}^*)$. Desired output is determined by expected demand $Y_{f,t+1}^* = Y_{f,t+1}^e$. The firm's information set in t consists of (i) the relative price $\frac{P_{f,t}}{P_t}$ – where $P_{f,t}$ is the price of the f -th good and P_t is the average price level – and (ii) excess demand

$$\Delta_{f,t} := Y_{f,t}^d - Y_{f,t} \quad (11)$$

where $Y_{f,t}^d$ is the demand for the f -th good and $Y_{f,t}$ is actual output. $\Delta_{f,t}$ shows up as a queue of unsatisfied customers if positive; as an inventory of unsold goods if negative. By assumption C-goods are not storable. Therefore involuntary inventories cannot be employed to satisfy future demand.

The firm makes use of two *rules of thumb* which govern price changes and quantity changes respectively.

The *price adjustment* rule is:

$$P_{f,t+1} = \begin{cases} P_{f,t}(1 + \mathbf{1}_u \rho_p) & \text{if } \Delta_{f,t} > 0 \\ P_{f,t}(1 - \mathbf{1}_o \rho_p) & \text{if } \Delta_{f,t} \leq 0 \end{cases} \quad (12)$$

where ρ_p is a random positive number, $\rho_p \sim \mathcal{U}(0, \bar{\rho}_p)$. $\mathbf{1}_u$ is an indicator function which takes value equal to 1 if the firm has underpriced the good (i.e., if $\frac{P_{f,t}}{P_t} < 1$), 0 otherwise. Analogously $\mathbf{1}_o$ is equal to 1 if the firm has overpriced (i.e., if $\frac{P_{f,t}}{P_t} > 1$), 0 otherwise.

Excess demand $\Delta_{f,t}$ and the relative price $\frac{P_{f,t}}{P_t}$ dictate the *direction* of price adjustment: the firm will increase (reduce) the price next period if it has registered excess demand (supply) and has underpriced (overpriced) the good in the current period. The *magnitude* of the adjustment is stochastic. The upper bound of the support of ρ_p limits the admissible price change. We also assume that the firm will never set a price lower than the average cost.

Since the quantity to be produced is equal to expected demand, the *quantity adjustment* rule takes the form of an updating rule for expected demand:

$$Y_{f,t+1}^* = Y_{f,t+1}^e = \begin{cases} Y_{f,t} + \rho_q \Delta_{f,t} & \text{if } \Delta_{f,t} > 0 \\ Y_{f,t} + \rho_q \Delta_{f,t} & \text{if } \Delta_{f,t} \leq 0 \end{cases} \quad (13)$$

where $\rho_q \in (0, 1)$. Both the direction and the magnitude of quantity adjustment are determined by excess demand.

Technology is represented by a Leontief production function, giving the maximum output the firm can produce in t : $\hat{Y}_{f,t} = \min(\alpha_N N_{f,t}, \alpha_K K_{f,t})$ where α_N and α_K represent labor and capital productivity respectively. Once a decision has been taken on desired output in $t + 1$, the firm retrieves from the production function how much capital it needs in $t+1$ to reach that level of activity (capital requirement): $K_{f,t+1}^* = Y_{f,t+1}^* / \alpha_K$. If actual capital $K_{f,t+1}$ is greater than the capital requirement, the desired *rate of capacity utilization* $x_{f,t+1} = \frac{K_{f,t+1}^*}{K_{f,t+1}}$ will be smaller than one. The labour requirement will be: $N_{f,t+1}^* = \frac{\alpha_K}{\alpha_N} K_{f,t+1}^*$. If actual employment in t , $N_{f,t}$, is smaller than the labor requirement in $t + 1$, the firm will post vacancies. If the opposite holds true the firm will fire workers. In this scenario, provided it is able to hire any additional required labor, the firm can reach the desired level of production.

On the other hand, if actual capital is smaller than the capital requirement, the former will be utilized at full capacity (the rate of capacity utilization will be $x_{f,t+1} = 1$ but desired output will not be reached: $Y_{f,t+1} = \alpha_K K_{f,t} < Y_{f,t+1}^*$).

Given a stock of undepreciated capital, actual capital in $t + 1$ $K_{f,t+1}$ will be determined by investment carried out in t , $I_{f,t}$. By assumption, in planning investment, the firm sets a *benchmark* equal to the capital stock used in production “on average” since the beginning of activity $\bar{K}_{f,t}$. This, in turn, is computed by means of an adaptive algorithm, i.e., the weighted average of past utilized capital from the beginning of activity until t with exponentially decreasing weights. In computing this weighted average, the firm employs a memory parameter $\xi_K \in (0, 1)$. Capital depreciates at the rate δ . Moreover we assume that C-firms may invest in each period with a probability π^k . Hence investment necessary “on average” to replace worn out capital is $\frac{\delta}{\pi^k} \bar{K}_{f,t+1}$.

We assume, moreover, that the firm plans to maintain, in the long run, a capital stock buffer. Therefore the *target* capital stock is equal to $K_{f,t+1}^T = \frac{1}{\bar{x}} \bar{K}_{f,t}$ where $\bar{x} \in (0, 1)$ is

the desired long run capacity utilization rate. Net investment is $K_{f,t+1}^T - K_{f,t}$. Therefore gross investment in t is:

$$I_{f,t} = \left(\frac{1}{\bar{x}} + \frac{\delta}{\pi^k} \right) \bar{K}_{f,t} - K_{f,t} \quad (14)$$

Once investment has been determined, C-firms go to the market for K-goods. The market protocol for this market follows the same lines as those of B-goods and L-goods. The f -th C-firm, with $f \in (1, N_F^c)$, has one go-to K-supplier (which sets the price P_0) and one potential new partner (which sets the price P_1) in the population of K-firms, indexed with $f \in (N_F^c + 1, N_F)$. If $R_{0,1}^k = \frac{P_0}{P_1} \leq 1$ the C-firm will stick to the go-to supplier and shop at the new partner only if the investment budget $I_{f,t}$ is not completely exhausted with the first purchase. It will switch to the new partner (and reverse the order of purchase) with probability $\pi^c = 1 - \exp \gamma_p(1 - R_{0,1}^k)$ if $R_{0,1}^k > 1$. If the C-firm's demand for K-goods has not been completely satisfied, it is forced to "save" the unspent portion of the investment budget. Therefore actual investment may turn out to be lower than planned investment.

The uniform nominal wage is set on the basis of labour market conditions captured by the distance between the current unemployment rate u_t and a threshold u^T . Whenever the unemployment rate is above (below) the threshold the wage will decrease (increase). The wage updating mechanism therefore is:

$$w_{t+1} = \begin{cases} w_t [1 + u_{up} (u^T - u_t)]; & u_t - u^T > 0 \\ w_t [1 + u_{down} (u^T - u_t)]; & u_t - u^T < 0 \end{cases} \quad (15)$$

where u_{up} and u_{down} are positive parameters. We will assume that $u_{up} > u_{down}$ to capture the downward stickiness of nominal wages.

A.2.2 K-firms

Firms indexed with $f \in (N_F^c + 1, N_F)$ are capital goods producers. The price adjustment rule followed by the f -th K-firm is equation (12) but the indicator functions should be re-interpreted. Denoting with $P_{f,t}$ the individual price and with P_t^k the average price of capital goods, the function $\mathbf{1}_o$ is equal to 1 if the K-firm in question has overpriced the good (i.e., if $\frac{P_{f,t}}{P_t^k} > 1$), 0 otherwise. Analogously, $\mathbf{1}_u$ takes value equal to 1 if the K-firm has underpriced the good (i.e., if $\frac{P_{f,t}}{P_t^k} < 1$), 0 otherwise.

The quantity adjustment rule departs from the one adopted by C-firms (see equation (13)) to take into account the fact that K-goods are durable and therefore storable: inventories of capital goods can be carried on from one period to another, depreciating

at a rate given by δ^k in each period. The quantity adjustment rule of the f -th K-firm is:

$$Y_{f,t+1}^* = \begin{cases} Y_{f,t} + \rho_q \Delta_{f,t} - Y_{f,t}^k & \text{if } \Delta_{f,t} > 0 \\ Y_{f,t} + \rho_q \Delta_{f,t} - Y_{f,t}^k & \text{if } \Delta_{f,t} < 0 \end{cases} \quad (16)$$

where $Y_{f,t+1}^*$ is the desired scale of activity, $Y_{f,t} + \rho_q \Delta_{f,t} = Y_{f,t+1}^e$ is expected demand, $Y_{f,t}^k$ is the inventory of firm f and $\Delta_{f,t}$ is excess demand for the K-good produced by firm f at time t . K-firms are endowed with a linear production function with labour as the only input.

A.3 The banking system

Once the quantities to be produced have been set and the cost of inputs determined, firms have to deal with financing. Consider a generic firm f . If the firm's internal liquidity (i.e., the current deposits held at the bank) $D_{f,t}$ are greater than the costs to be incurred, the firm can finance production and investment (if any) internally. If, on the other hand, liquidity is not sufficient to carry out production and investment up to the desired level, the firm applies for a loan to fill its financing gap which is given by

$$F_{f,t} = wN_{f,t} + \mathbf{1}_c P_{t-1}^k I_{f,t} - D_{f,t} \quad (17)$$

where $\mathbf{1}_c$ is an indicator function which assigns value 1 to C-firms and 0 to K-firms (since only C-firms purchase capital goods). We assume that C-firms assess the financing gap (and the demand for loans) before accessing the market for capital goods. Hence capital goods to be bought in t are priced with the ‘‘average’’ price of capital goods P_{t-1}^k .

For simplicity we assume there is only one bank which collects deposits from firms and households, supplies credit to firms and purchases government bonds. The bank decides (i) the interest rate to be charged to each borrower and (ii) the size of the loan (which may be different from the borrower's financing gap). Both decisions are affected by the borrower's leverage $\lambda_{f,t}$:

$$\lambda_{f,t} = \frac{L_{f,t}}{E_{f,t} + L_{f,t}} \quad (18)$$

where $L_{f,t}$ is the firm's debt and $E_{f,t}$ is equity or net worth.

The bank makes an assessment of the probability of default, which is increasing with leverage. The perceived probability of default for the f -th C-firm, $f \in (1, N_F^c)$, is:

$$\pi_{f,t}^b = \frac{e^{b_{0c} + b_{1c} \lambda_{f,t}}}{1 + e^{b_{0c} + b_{1c} \lambda_{f,t}}} \quad (19)$$

with $b_{0c} < 0$ and $b_{1c} > 0$. Analogously, the assessed probability of default for the f -th

K-firm, $f \in (N_F^c + 1, N_F)$, is:

$$\pi_{f,t}^b = \frac{e^{b_{0k} + b_{1k}\lambda_{f,t}}}{1 + e^{b_{0k} + b_{1k}\lambda_{f,t}}} \quad (20)$$

with $b_{0k} < 0$ and $b_{1k} > 0$. The interest rate charged by the bank to each firm is determined by adding an *external finance premium* (Bernanke et al., 1996) to the exogenous risk free interest rate r . The external finance premium is increasing with the probability of default which in turn is (non-linearly) increasing with leverage. The interest rate charged to the generic f -th firm is:

$$r_{f,t} = \mu f(r, \lambda_{f,t}) \quad (21)$$

where the function $f(\cdot)$ is increasing in both arguments.²⁶

In order to determine the size of the loan given to a firm f , the bank first sets a tolerance level for the potential loss $\Gamma_{b,t}$ on credit extended to any individual borrower as a fraction ϕ_b of its net worth: $\Gamma_{b,t} = \phi_b E_{b,t}$. The borrower's total debt in t will be $\Phi_{f,t} + L_{f,t-1}$ where $\Phi_{f,t}$ is the new credit line to be supplied in t . We assume the bank sets the new credit line in order to equate the expected loss on loans extended to the f -th firm to the tolerance level: $(\Phi_{f,t} + L_{f,t-1})\pi_{f,t}^b = \phi_b E_{b,t}$. Therefore the new credit line is:

$$\Phi_{f,t} = \frac{\phi_b}{\pi_{f,t}^b} E_{b,t} - L_{f,t-1} \quad (22)$$

Given the current exposure of the bank to the firm, the new credit line is increasing with the bank's net worth and decreasing with the firm's leverage. The size of the loan actually granted to firm f at time t will be the minimum between the new credit line and the financing gap:

$$\dot{L}_{f,t} = \min(\Phi_{f,t}, F_{f,t}) \quad (23)$$

If the latter is greater than the former the firm will be rationed on the credit market and therefore forced to scale down its investment and/or production. In addition to making interest payments, firms in each period repay a fraction ζ of their total debt to the bank. The bank remunerates deposits and earns interests on loans and on Government bonds. The interest rate on deposits is determined by marking down the risk-free interest rate.

A.4 Net worth updating

In every period, each firm's net worth E_f is updated by means of retained net profits:

$$E_{f,t+1} = E_{f,t} + (1 - t_\Pi)(1 - \omega)\Pi_{f,t} \quad (24)$$

²⁶For the specification of $f(\cdot)$ see Assenza et al. (2015).

Also the bank's net worth is updated by means of retained profits:

$$E_{b,t+1} = E_{b,t} + (1 - t_{\Pi})(1 - \omega)\Pi_{b,t} - BD_t \quad (25)$$

where $\Pi_{b,t}$ is the bank's profit and BD_t is *bad debt*, i.e., the book value of non-performing loans.

A.5 Entry-exit mechanism

If the liabilities of a firm exceed its assets (so that its equity turns negative), it is assumed to go bankrupt and exit.²⁷

As to entry, a newly born firm will enter sector $j = B, L, K$ and replace a bankrupt firm in the same sector with probability $\pi_{j,t}^e = [1 + \exp(\gamma_e \kappa_t^j)]^{-1}$, which increases with the average profit rate prevailing in the sector κ_t^j . The firm owner of the previously bankrupt firm being replaced will provide the initial equity injection to the entrant firm. In a sense, a firm which goes bankrupt is 'dormant', remaining inactive for a variable number of periods until a new firm succeeds in entering and replacing it. At any given time, therefore, the number of *active* firms may be smaller than N_F . N_F itself is fixed and the number of active firms cannot exceed it.

Regarding the bank, we assume that if its equity becomes negative due to losses on bad debt, a bail-in procedure will immediately be applied: all firm-owners (who collectively own the bank) make a transfer to the bank to make its equity positive again.

A.6 The Public Sector

The public sector taxes wages and profits, provides unemployment subsidies, sick-pay and pensions (to workers), makes interest payments on government bonds (to the bank) and carries out government expenditure on healthcare. The latter is a constant fraction g of full employment output, taking the initial population of active workers N_A as a basis for calculation. In symbols:

$$G = g\alpha_N N_W \quad (26)$$

We assume this expenditure translates one for one into a supply of healthcare services to the population via the public healthcare system. G is in the first instance allocated to firms according to their relative revenue in the previous period.²⁸

²⁷The firm can also be illiquid. If a firm's liquidity (bank deposit) is smaller than zero at the end of the period but its equity is positive, it receives a transfer from the firm owner to make up the negative balance. If, after the transfer, the firm's liquidity is still negative, the bank takes a loss equal to the negative balance and the firm's deposit become zero. However, the illiquid firm does not exit the economy unless its equity also turns negative.

²⁸The f -th firm receives demand from the public sector equal to the fraction $\frac{R_f}{\sum_{f=1}^{N_F} R_f}$ where R_f represents the firm's revenue. If, after the first round of expenditure, the government has been unable to

A public sector deficit occurs when taxes turn out to be lower than the sum of subsidies, interest payments and government expenditure on healthcare. In this case, the government will issue new bonds. For simplicity, we assume that all government bonds are purchased by the bank at the fixed risk-free rate r .

A.7 Demand and supply of healthcare

We assume that economic decisions are taken every month while the health component of the model runs at a weekly frequency. We will indicate the current week with the subscript τ .

First we define a benchmark case, the **Normal Times (NT)** scenario. We assume that during normal times, in any given week, a healthy agent may catch a non-infectious disease with a certain probability π^i . The presence of this disease in turn generates a demand for healthcare services. The non-infectious disease is also not lethal: after a fixed number D_d^n of weeks²⁹ the sick recover. Recovery does not imply immunity: recovered agents may randomly become susceptible again in the future. This assumption implies that the non-infectious disease is endemic.

For simplicity we assume that only workers (both active and inactive) may get ill with the non-infectious or the epidemic disease (described below). Since, as outlined below, age is an important factor affecting the course of the epidemic disease, we divide the population of workers into three age-segments. We denote with ϕ_y , ϕ_m and ϕ_o the fractions of the population consisting of young, middle-aged and old workers.³⁰ The variable age_h assumes value 1 if the agent is young – i.e., if they belong to the fraction ϕ_y of the population – 2 if middle-aged and 3 if old. The h -th sick agent, $h \in (1, N_W)$, generates a demand of healthcare in week τ – denoted with $H_{h,\tau}^d$ which is increasing with age and affected by an idiosyncratic shock:

$$H_{h,\tau}^d = h_1 age_h + h_2 u_{h,\tau} \quad (27)$$

where $u_h \sim \mathcal{U}(0, 1)$.

The total supply of healthcare services in every period τ is given by G . In the first instance, this is allocated to agents who were already ill and receiving healthcare in the previous period and who still require it. The residual is then allocated to patients who have fallen ill in the current period: A randomised queue of all agents requiring and not already receiving healthcare is formed and agents are admitted into the healthcare

spend the entire amount G (because some of the firms did not produce sufficient output), the remaining demand is redistributed between those firms which still have goods available until the exact amount G has been spent.

²⁹ $D_d^n = 4$ in the simulations shown below.

³⁰In our calibration, $\phi_y = 0.15$, $\phi_m = 0.65$ and $\phi_o = 0.2$. These parameter values roughly capture the current composition by age of the population of Lombardy.

system until the residual supply is exhausted. Hence the demand for healthcare may be rationed. If an agent's demand exceeds the remaining supply of healthcare, that agent is rationed and receives only a fraction of the desired supply. All subsequent agents in the queue are rationed completely. All rationed agents will queue again in the next period if they still require care.

Sick agents who were previously in the labour force become inactive and receive sick-pay. Old people are inactive by assumption and receive pensions. Retired agents who become sick will continue to receive the pension instead of sick-pay.

Appendix B The epidemiological sub-model

In this section we describe the dynamics of an epidemic, i.e., an *infectious* disease.³¹ The epidemic differs from the non-infectious disease because of the transmission from one subject to the others through *contagion*. Despite being based on contagion through a network structure of contacts between agents, the epidemiological sub-model is similar to classic compartmental models in that agents can be classified into various states. The notation is as follows:

\mathcal{I}_τ denotes the *cumulative* number of (both detected and undetected) infections from the beginning of the epidemic up to period τ . $\dot{\mathcal{I}}_\tau$ denotes the number of *new* infections in τ . $\mathcal{I}_{c,\tau}$ denotes the stock of *currently* infected agents in τ . $\Delta\mathcal{I}_{c,\tau}$ denotes the *change* in the stock of currently infected in τ . Note that in general, $\Delta\mathcal{I}_{c,\tau} \neq \dot{\mathcal{I}}_\tau$, since the former includes newly recovered agents (and hence may be negative) whilst the latter only includes new infections and hence must be ≥ 0 . Similarly, let \mathcal{D}_τ denote the cumulative number of *detected* infections, with $\dot{\mathcal{D}}_\tau$, $\mathcal{D}_{c,\tau}$ and $\Delta\mathcal{D}_{c,\tau}$ having the equivalent interpretations of the variables defined above. \mathcal{M}_τ is the cumulative number of deaths, \mathcal{H}_τ the cumulative number of agents requiring healthcare due to the epidemic disease, and \mathcal{R}_τ the cumulative number of recoveries from the epidemic disease. For all three, we also define the respective derivative variables as above.³² Finally, $\mathcal{S}_{c,\tau}$ denotes the stock of agents who are currently susceptible to the disease and $\Delta\mathcal{S}_{c,\tau}$ the change in this stock.³³

The epidemic begins in an exogenously determined week labelled τ_E , in which a small number of workers are exogenously infected with the epidemic disease. These people are the *initial infected* (and infectious) and will be denoted with \mathcal{I}_{c,τ_E} .³⁴ The (healthy and)

³¹In the following we will not use the notion of a pandemic because we model a closed economy and there is no diffusion of the epidemic across borders.

³²Since agents in the model can be infected multiple times with the epidemic disease, provided that they recover in-between, and given that dead agents can be replaced by newly born ones, it makes sense to define the stocks of currently recovered and dead agents as distinct from the respective cumulative values.

³³Obviously, it makes little sense to define a cumulative stock of susceptible agents or a change therein.

³⁴In the simulations, we set $\mathcal{I}_{c,\tau_E} = 5$. The fraction of the infected in the initial population therefore

susceptible agents after the appearance of the infected are $\mathcal{S}_{c,\tau_E} = N_W - \mathcal{I}_{c,\tau_E}$ since at the beginning of the epidemic, all N_W workers in the model are alive. These susceptibles may then be infected by the initial infected in τ_E and subsequent periods as described below.

Some infected agents develop mild symptoms or do not develop symptoms at all (*non-symptomatic* for short). In this case the infection can be detected only if the agent is subjected to a test. In each period, every undetected infected agent may be detected with a probability π_τ^r which, as explained in the main text, is fixed in the uncontained epidemic but becomes endogenous in the social distancing and lockdown scenarios. Agents who test positive are quarantined and therefore cannot spread the disease. Undetected infected can still spread the disease. People who develop serious symptoms are detected with certainty. \mathcal{D}_τ , the cumulative number of detected infections, hence includes all infections leading to serious symptoms as well as all infections detected through tests on non-symptomatic agents. The probability for an agent to develop serious symptoms is increasing with age.³⁵ All agents developing serious symptoms require healthcare and hence become part of \mathcal{H}_τ . Their individual demand for healthcare is given by equation (27). All agents who are currently infected and detected, $\mathcal{D}_{c,\tau}$ will be inactive (and receive sick pay if they are not retired) and will not have contacts with other agents for the entire duration of the disease. For simplicity, we assume that the infected remain contagious for the entire duration of the illness. The undetected infected therefore can spread the disease for the entire duration of their illness.

Non-symptomatic agents recover with certainty after a certain number of periods of being ill.³⁶ Agents developing serious symptoms, on the other hand, may die with some probability during each period of the illness before recovering. In period τ , the h -th infected agent with serious symptoms will face a probability of death which is increasing with age and with excess demand for health care:

$$\pi_{h,\tau}^m = \widehat{\pi}_\tau^m age_h^3 + h_{3,\tau}(H_{h,\tau}^d - H_{h,\tau}^s) \quad (28)$$

where $\widehat{\pi}_\tau^m > 0$ and $h_{3,\tau} > 0$, $H_{h,\tau}^d$ is the agent's demand for health care and $H_{h,\tau}^s$ is the amount of healthcare they actually receive, which depends on the free capacity of the healthcare system. We assume that both $\widehat{\pi}_\tau^m$ and $h_{3,\tau}$ decrease over time even in the absence of a vaccine until they reach a lower bound. A rationale for this assumption is that even without a vaccine, healthcare systems may over time become better at treating

is $5/30000 = 1/6000$. Since Lombardy has a population of approximately 10 million, this means that the set of the initial infected (which, it should be noted, are all initially *undetected*) in the model is equivalent to $\frac{10^7}{6000} = 1666$ people in Lombardy.

³⁵In the simulations we assume that this probability is $\pi_y^h = 0.01$ for the young, $\pi_m^h = 0.02$ for the middle aged and $\pi_o^h = 0.525$ for the old.

³⁶In the present calibration, the duration of the infectious disease – denoted with d^i – is drawn from a uniform distribution over the interval $3 \leq d^i \leq 5$ for any infected individual.

a novel disease (in the case of Covid-19 this may involve the use of existing medicines, increasing experience as to when patients should be intubated, etc.). The laws of motion are:

$$\begin{aligned}\widehat{\pi}_\tau^m &= \max(\underline{\pi}^m, \widehat{\pi}_{\tau-1}^m(1-z)) \\ h_{3,\tau} &= \max(\underline{h}_3, h_{3,\tau-1}(1-z))\end{aligned}\tag{29}$$

Instead of postulating the law of motion of the number of infected people as in SIR models, we adopt a granular approach to contagion focusing on networks in order to depict the transmission of the epidemic among agents.

Contagion spreads in three networks: the workplace (employment network), the marketplace (shopping network) and social relations (social network). Employed workers are nodes in the *employment network*. Each employed worker is linked to all co-workers in the firm they work for meaning that they encounter them every week. If a firm goes into smart working, only a share of possible workplace encounters take place. If a firm is shut down by a lockdown, no workplace encounters occur at that firm.

In addition, all worker households are nodes in the *shopping network*. A certain number of households shop at a given C-firm. If one of these buyers is infective, they can spread the disease to other households shopping at the same firm. We list all possible connections between the customers of a given firm and assume that a fixed share ($\frac{1}{3}$) of those encounters actually take place (reflecting the assumption that not all customers visit the firm at the same time). This share is reduced if there is a lockdown in place or people engage in social distancing.

Finally, we build a *social network* to depict encounters during leisure time. Each worker household has a set of social connections consisting of family and friends. The total number of social connections is a (very small) fraction of the maximum number of possible undirected connections between worker households, $\frac{N_w(N_w-1)}{2}$.³⁷

While both the employment and shopping networks change over time as households change employment and the firms at which they shop, the social network is static.

We assume that each infected and undetected agent meets all the agents they are connected to (at work, while shopping and during leisure time) in every week. Let N_τ^C denote the number of connections in week τ which involve exactly one undetected infected and one susceptible agent. We assume that only a fraction (the transmission rate) of these connections may lead to a new infection. In other words, there is a *maximum number* of potential *new* infections in week τ given by

$$\dot{I}_\tau^{max} = \rho_{c,\tau} N_\tau^C\tag{30}$$

³⁷In the present calibration, the total number of connections in the social network is $\frac{1}{15000}$ of all possible connections.

where $\rho_{c,\tau} \in (0, 1)$ is the transmission rate which incorporates a seasonal effect, being lower from May to September and higher from October to April. We then take a sample of size \dot{I}_τ^{max} from the set of connections involving exactly one undetected infected and one susceptible agent. The sample is weighted such that the likelihood of being drawn is highest for social connections, second-highest for employment connections and lowest for shopping connections. In the absence of social distancing, each of these connections leads to an infection with certainty. Under social distancing, a connection leads to an infection with probability $1 - \beta$ if one agent involved in the connection is socially distancing and with probability $1 - 2\beta$ if both agents are socially distancing.

As indicated above, the infection with the epidemic disease ends either with recovery or death. If an infected agent recovers, the stock of currently recovered agents, $\mathcal{R}_{c,\tau}$ increases by one. If the agent was previously economically active and became inactive due to their infection being detected, they will re-enter the labour force as an unemployed agent and begin to look for a job. If an infected agent dies, the stock of current dead, $\mathcal{M}_{c,\tau}$ increases by one. We assume that there are no bequests, such that the assets of dead agents are simply written off. We assume that recovered agents may become susceptible to the disease again and that dead agents can be replaced by newly born ones. At the end of every week τ , each agent in the stock of currently recovered may become susceptible to the epidemic disease again with a low probability π_e^s . An agent may hence be infected with the epidemic disease more than once over the course of a simulation run. This factor makes it more likely for the disease to become endemic. At the beginning of every month t , every dead agent may be replaced with a new agent of age 1 (i.e. a young agent) with a low probability π^n ($= 0.0125$ in the simulations shown here). If the dead agent is replaced, the stock of currently dead agents decreases by one. This replacement mechanism ensures that when the disease becomes endemic, it does not cause the entire population to eventually die out. Since the probability of death never goes to zero, the absence of a replacement mechanism for dead agents would mean that unless the disease dies out first, the model population would go to zero in the very long run (i.e. beyond the time horizons simulated in this paper).

Appendix C Parameter values

Tables 5 and 6 below provide the lists of model parameters pertaining to the macroeconomic sub-model and the epidemiological sub-model respectively.

Table 5: Macroeconomic sub-model parameters

Symbol	Description	Value
N_W	Number of workers	30000
N_F^b	Number of B-firms	360
N_F^l	Number of L-firms	240
N_F^k	Number of K-firms	200
z_e	Number of Firms visited by unemployed	5
ξ_Y	Memory parameter for human wealth	0.55
c_W	Propensity to consume out of financial wealth	0.00835
ρ_q	Quantity adjustment parameter	0.2
$\frac{\rho_p}{p}$	Price adjustment upper bound	0.08
μ	Bank's gross mark-up	1.007
δ	Capital depreciation rate	0.01
π^k	Probability to invest	0.4
ϕ_b	Bank's leverage parameter	0.0025
ζ	Debt repayment rate	0.05
ξ_K	Memory parameter for capacity utilisation	0.2
α_N	Labour productivity	2/3
α_K	Capital productivity	1/6
ω	Dividend payout ratio	0.25
\bar{x}	Target capacity utilisation	0.85
δ^k	Inventory depreciation	0.08
b_{0c}	Bank's risk evaluation parameter (C-firms)	-10
b_{1c}	Bank's risk evaluation parameter (C-firms)	10
b_{0k}	Bank's risk evaluation parameter (K-firms)	-15
b_{1k}	Bank's risk evaluation parameter (K-firms)	15
r	Risk-free interest rate	$\frac{0.01}{3}$
r_d	Interest rate on deposits	$\frac{r}{2}$
s_u	Replacement rate (unemployment subsidy)	0.75
s_p	Replacement rate (pension)	0.9
s_s	Replacement rate (sick-pay)	0.75
t_w	Tax rate on wage income	0.275
t_Π	Tax rate on profits	0.3
u^{up}	Upward wage adjustment parameter	$\frac{0.1}{3}$
u^{down}	Downward wage adjustment parameter	$\frac{0.01}{3}$
u^T	Unemployment threshold	0.1
g	Ratio of government healthcare expenditure to full employment GDP	0.04
γ_p	Probability of switching parameter	40
γ_e	Probability of entry parameter	-40

Table 6: Epidemiological model parameters

Symbol	Description	Value
ϕ_y	Share of young agents in the population	0.15
ϕ_m	Share of middle-aged agents	0.65
ϕ_o	Share of old agents	0.2
π^i	Probability of catching the normal disease	0.0012
D_d^n	Duration normal disease	4
	Susceptibility probability normal disease	0.1
σ^L	Consumption shock to L-goods parameter (baseline and lower bound)	[1.65e-3,1.65e-4]
σ^B	Consumption shock to B-goods parameter (baseline and lower bound)	[5.5e-4,5.5e-5]
π_y^h	Share of young agents with serious symptoms	0.01
π_m^h	Share of middle aged agents with serious symptoms	0.02
π_o^h	Share of old agents with serious symptoms	0.525
	Total number of possible connections	449985000
	Number of permanent connections	29999
	Share of deactivated L-firms in lockdown	1/3
d_{max}^{lock}	Lockdown maximum duration (weeks)	12
\dot{D}_{lock}	Lockdown activation threshold (new detected)	30
\dot{D}_{end}	Lockdown lifting threshold	12.5
d^i	Duration of epidemic disease (weeks)	$\ \mathcal{U}(3, 5)\ $
z	Post-lockdown adjustment parameter	0.0775
	Share of connections under lockdown	0.25
	Share of work connections under lockdown	0.375
	Share of shop connections (out of weekly visitors)	1/3
	Share of shop connections under lockdown	(1/3)·0.25
c_{SD}	Cost of distancing (in lockdown)	6 (-6)
ι	Persistence of distancing index	0.725
β	Distancing effect on infection probability	1/3
$\rho_{c,\tau}$	Transmission rate (October to April)	0.07
$\rho_{c,\tau}$	Transmission rate (May to September)	0.04
$\overline{D}_{c,SD}$	Social distancing threshold	5
\mathbf{N}	Intensity of choice for social distancing	[0.05, 0.5, 1]
h_1	Health demand parameter	2
h_2	Health demand parameter	0.1
h_3	Death probability parameter (baseline and lower bound)	[0.0375,0.0125]
π^m	Death probability (baseline and lower bound)	[0.0075,0.0025]
π^r	Detection probability (baseline and upper bound)	[0.02,0.125]
γ_d	Adjustment of detection probability	0.0005
π_e^s	Susceptibility probability epidemic disease	0.0025
\mathcal{I}_0	Number of initially infected	5
π^n	Replacement probability of dead agents	0.0125

Table 7: Vaccine parameter

Description	Value
Coverage rate (initial value)	0.01
Weekly increments and upper bound of coverage rate	[0.001,0.05]
Vaccine efficacy w.r.t. contagion	0.70
Vaccine efficacy w.r.t. serious disease	0.90
Vaccine-induced immunity duration (weeks)	$\mathcal{N}(52, 2)$

Table 8: Variants parameter

Description	Value
Number of initially infected with variant	5
Increase factor transmission rate	1.5
Reduction factor distancing effect	0.25
Variant 1,2: reduction factor vaccine efficacy w.r.t contagion	[0.80,0.80]
Variant 1,2: reduction factor vaccine efficacy w.r.t. serious disease	[1,0.80]