ABC: An Agent Based Exploration of the Macroeconomic Effects of Covid-19

Domenico Delli Gatti, Severin Reissl

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>

ABC: An Agent Based Exploration of the Macroeconomic Effects of Covid-19

Abstract

We employ a new macro-epidemiological agent based model to evaluate the “lives vs livelihoods” trade-off brought to the fore by Covid-19. The disease spreads across the networks of agents’ social and economic contacts and feeds back on the economic dimension of the model through various channels such as employment and consumption demand. We show that under a lockdown scenario the model is able to closely reproduce the epidemiological dynamics of the first wave of the coronavirus epidemic in Lombardy. We then explore the efficacy of the fiscal response to Covid-19 which may take different routes: income support, liquidity provision, credit guarantees. In an agent based setting we gain additional insights on the way in which fiscal measures impact not only on GDP but also on the defaults of firms and the allocation of inputs. We find that liquidity support for firms, a short-time working scheme with compensation for workers, and direct transfer payments to households are effective policy tools to alleviate the economic impact of the epidemic and the lockdown.

JEL-Codes: E210, E220, E240, E270, E620, E650.

Keywords: agent-based models, epidemic, Covid, fiscal policy.

*Domenico Delli Gatti**
Department of Economics and Finance
Catholic University Milan / Italy
domenico.delligatti@unicatt.it

Severin Reissl
IUSS Pavia / Italy
severin.reissl@iusspavia.it

*corresponding author

December 5, 2020

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 721846 (ExSIDE ITN).

1 Introduction

The impact of Covid-19 on every aspect of economic life is at the centre of the interest of the profession since the outbreak of the epidemic. The literature has grown exponentially – becoming huge in a matter of weeks – in tandem with the progression of the contagion and does not show any tendency to level off.

A substantial part of this literature focuses on the intertwined health and macroeconomic effects of the epidemic. There is a clear pattern in the consensus approach to this issue: either a macroeconomic model is augmented with a health sub-model or an epidemiological model is augmented with a macroeconomic sub-model. Some authors place the emphasis on the macroeconomic component, some others on the epidemiological component but in both cases the end result is a *macro-epidemiological model* which captures the two-way feedback between the health condition of the population on one hand and the macroeconomy on the other. Such models are then put to work to assess the combined macroeconomic and health effects of (public health) *mitigation policies* – primarily the lockdown – and of (macroeconomic) *stabilization policies*, mainly demand management. This modelling strategy is certainly fruitful – as shown by the proliferation of papers and models – and has provided important insights into the pervasive effects of Covid-19 and the relative efficacy of mitigation and stabilization policy. While the epidemiological sub-models are typically variants of the SIR framework, the macroeconomic sub-models spring from different theoretical approaches, so that architectures and behavioural rules differ from one model to the other. So far, however, most of the macro-epidemiological models are characterized by canonical (new classical or new Keynesian) microeconomic foundations and design of the macroeconomic structure.

In time of Covid, however, the granularity and flexibility of the agent based approach to modelling the macroeconomy can be particularly useful. As Paul Gourinchas put it “A modern economy is a complex web of interconnected parties: employees, firms, suppliers, consumers, banks and financial intermediaries... Everyone is someone else’s employee, customer, lender, etc. A sudden stop [such as Covid-19] can easily trigger a cascading chain of events, fueled by individually rational, but collectively catastrophic, decisions.”

(Gourinchas, 2020, p. 33).

Exciting new research to assess the economic consequences of Covid-19 using a granular approach is in full swing, in particular within the GROWINPRO project funded by the European Commission, at Eurace@Unibi, at INET-Oxford and at IIASA. There are very few attempts, however, to provide an integrated *agent-based macro-epidemiological model*. This is surely a gap in the literature.

We contribute to filling the gap by proposing a new agent based macro - epidemiological model of Covid-19 – hence the ABC label – based on the CATS framework presented by Assenza et al. (2015).¹ We adopt the strategy of integrating an agent based macroeconomic sub-model and an epidemiological sub-model. Our strategy is similar to that of Basurto et al. (2020) who integrate a SIR epidemiological component into an adapted Eurace@Unibi framework. Their model is fairly large in terms of the number of agents – the initial population consisting of 100 thousand households and around 4 thousand firms – so that some of the features of Eurace@Unibi have been shut off. We opted for a much smaller scale – ABC features 2800 households and 300 firms – to make the computational burden less heavy, allowing us to retain all features of the original CATS framework and add some suitable extensions.²

In both sub-models we exploit the *granularity* and *flexibility* of the agent-based approach. Granularity is a foundational feature of the macroeconomic component of ABC, which is a version of Assenza et al. (2015) extended and modified to enable the analysis of the economic effects of an epidemic disease.

We also apply the granular approach to the epidemiological component. The canonical SIR model is essentially a system of differential equations describing the laws of motions of aggregates, i.e., the epidemiological groups – Susceptibles, Infected, Removed (dead or recovered). In a sense, therefore, it is an aggregative model: it is not grounded in the actions of the individual agents but captures the dynamics of the aggregates. Our

¹See also Assenza et al. (2018) and Delli Gatti and Grazzini (2020).

²There is clearly a trade-off involved in this choice. A smaller population of agents allows us to consider a fairly rich macroeconomic and epidemiological environment. At the same time, as will become clear below, the reduced dimensionality – relative to Basurto et al. (2020) – makes it difficult to reproduce certain characteristics of the epidemic, particularly the (delayed) emergence of second waves following strict lockdowns during the first wave.

starting point in describing the transmission of contagion is the observation that the disease spreads through the agents’ interactions in social and economic networks. The evolution of the disease both influences and depends on the behaviour of individual agents along the ramifications of these networks. Hence, also in the epidemiological submodel of ABC we start from a population of heterogenous agents. We track agents’ interactions along the networks of their contacts both in the workplace and at home, i.e., during leisure time. The dynamics of each epidemiological group therefore can be reconstructed “from the bottom up” by summing across individuals who happen to be in the same health condition. The epidemiological sub-model of ABC, therefore can be characterized as a *network-based* SIR model.

While contagion may spread very quickly, making it important to depict epidemiological phenomena at a relatively high frequency, macroeconomic phenomena are typically modelled at lower frequencies. We exploit the flexibility of the agent-based approach which allows models to encompass processes that run at different time frequencies. In ABC, the basic unit of time of the epidemiological sub-model is one week and that of the macroeconomic submodel is one month.

Ours is not just a modelling exercise. We put the model to test to face two research issues:

- (A) First we want to assess the *trade off* (if any) between “saving lives” and stabilizing GDP – “lives vs livelihood” – by replicating epidemic curves and recession profiles with and without mitigation measures, primarily the lockdown. In an agent based setting we gain additional insights on the change of relative prices, defaults and the reallocation of inputs, as well as fiscal conditions (when only automatic stabilizers are endogenously activated).
- (B) Second we want to assess the relative *efficacy* of stabilization policies adopted to cope with the downturn created by the epidemic and exacerbated by the lockdown. We will focus on the fiscal response to Covid-19 which may take different routes: income support of one sort or another, liquidity provision, credit guarantees and so on. In an agent based setting we gain additional insights on the way in which

discretionary fiscal policies impact not only on GDP but also on the defaults and the allocation of inputs.

We tackle issue (A) by carrying out a number of experiments to assess the effect of *mitigation policies* – primarily the lockdown – on the trade off between saving lives and saving GDP.

To carry out this task we calibrate the model to reproduce moments drawn from macroeconomic data for the Lombardy region and use ABC to simulate alternative epidemiological scenarios. We characterize the baseline scenario as **Normal Times**. In normal times people get sick because they can be hit by a normal (non-infectious and non-lethal) disease. In this scenario, therefore, the epidemiological component of the model is shut off: (i) sick people do not transmit their disease to the others and (ii) they always recover (i.e., they do not die and total population does not change). The baseline simulations are characterized by irregular fluctuations around a long run mean (a quasi-steady state) of the most important macroeconomic variables. They replicate fairly closely the stylized macroeconomic facts of business fluctuations in Lombardy. In the baseline scenario, ABC also generates a “normal” demand for healthcare services and a “normal” level of health-motivated Government outlays (expenditure for the public provision of healthcare services and sick-pay transfers to households).

We then consider different epidemiological scenarios and compare the simulated time series under each scenario with the corresponding time series in normal times. In each scenario, people can catch an infectious and potentially lethal disease. Hence, over the “long run” – i.e., when the flow of newly infected people tends to zero and the stock of infected people reaches a plateau – a fraction of the population is permanently “removed” because of death.

First we simulate the **Uncontained Epidemic** scenario. It turns out that, if uncontained, the epidemic not only has a huge death toll but it also carries a significant output loss *over an extended time horizon*. The level of GDP goes down – albeit slowly – at the outbreak of the epidemic (short run) and stabilizes over time around a quasi-steady state permanently below the level of normal aggregate activity (long run).

The death toll is due to a large extent to the pressure of exponentially increasing demand for healthcare on the public health system. The macroeconomic loss (decline of GDP) is due to the supply and demand effects of the epidemic. Workers get sick and become inactive, employment goes down and firms cut production. Therefore, *even in the absence of mitigation measures*, the epidemic represents a non-negligible and lasting *negative supply shock*. In addition, there is also a *negative demand shock*: Firstly, people who get sick or become unemployed suffer a significant curtailment of their incomes (as they will receive sick-pay or unemployment benefits instead of wages) and their consumption contracts. Secondly, the consumption expenditure of people who die from the disease is permanently “removed” from the macroeconomy, hence aggregate consumption falls. If the model provides any indication, it suggests that just letting the epidemic run to avoid closures of economic activity does not benefit the macroeconomy (as it does not avoid a contraction of aggregate output) especially over the long run (as the economy stabilizes in a lower quasi-steady state).

The second scenario we consider is characterized by (endogenous) **Social Distancing**. As the label suggests, in this scenario the contagion is endogenously restrained by the (voluntary) social distancing behaviour that agents adopt when the epidemic gets “out of control”, i.e., when the number of infected people overcomes a “psychological” threshold. Voluntary social distancing substantially reduces the death toll. Moreover, the convergence to the plateau occurs in an oscillatory fashion, meaning that endogenous social distancing gives rise to multiple waves of the epidemic disease. After the first wave, as the number of infections becomes smaller people relax and cease to engage in social distancing but this relaxation boosts the number of infections. Here comes the second wave, which leads to infections overcoming the threshold again and makes people re-enter the social distancing mood. The macroeconomic loss in this case is bigger than in the uncontained epidemic case in the short run – immediately after the outbreak of the epidemic – but becomes smaller over a longer horizon as the death toll is considerably smaller.

Finally we explore the **Lockdown** scenario. In this case a relevant fraction of firms is shut down. The lockdown has a remarkable mitigating effect on the lethality of the

disease: the death toll is one order of magnitude smaller than in the uncontained epidemic case. In the short run the macroeconomic loss is dramatic but GDP bounces back over a longer horizon. The trade off between lives and livelihood is remarkable in the early stage of the lockdown but fades away in the long run. Having calibrated the model on the case of Lombardy, we show that ABC under lockdown is able to closely reproduce the infection and fatality numbers for Lombardy where the lockdown was activated fairly soon (first week of March). Using these simulations, we could infer that in the absence of a lockdown, the number of fatalities in Lombardy could have been 20 times bigger than the actual figure.

We tackle issue (B) by carrying out a number of policy experiments to assess the efficacy of macroeconomic *stabilization policies* in counteracting the epidemic- and lockdown-induced downturns. We consider different measures – ranging from income and employment support to liquidity provision and credit guarantees for firms – and in each experiment we compare the simulated time series generated by a specific measure with the corresponding time series under the lockdown scenario. In the simulations, the most effective policy measures appear to be income support for households as well as a combination of employment support for workers (in the form of a ban on layoffs and a redundancy fund) and temporary liquidity support for firms. These policies significantly speed up recovery and strongly reduce the output loss relative to the lockdown scenario without any stabilisation policy.

The paper is structured as follows. After a brief review of the literature (section 2), sections 3 to 5 provide a detailed description of the ABC model, including the epidemiological component. Section 6 discusses the calibration and validation of the baseline scenario (normal times). Section 7 presents the epidemic curves under different epidemiological scenarios while section 8 focuses on the macroeconomic effects of these scenarios. Macroeconomic stabilization measures to counter the fallout from the epidemic and the lockdown are discussed in section 9. Section 10 concludes.

2 A concise review of the literature

The literature on the effects of Covid-19 and on the policies to contain the epidemic and reduce the economic damage is already huge. In this section we focus only on the macroeconomic consequences of Covid-19, and only on some of the more significant contributions in this field.

The descriptive literature and the public debate³ have revolved around two major questions:

1. What kind of shocks hit the macroeconomy following the outbreak of an epidemic?
What are the channels of transmission of these shocks?
2. Is there a trade off between “lives” and “livelihoods”? Do public health measures aimed at containing the death toll have small and rapidly disappearing or sizable and prolonged effects on aggregate economic activity?

As to the first question, a consensus has rapidly emerged: the epidemic (and the associated public health containment measures) generates both a demand and a supply shock with profound contractionary consequences (Baldwin and Weder di Mauro (2020a)). The most intriguing characterization of this double shock has been put forward, in our view, by Guerrieri et al. (2020), who characterize Covid-19 as a *Keynesian supply shock*, i.e. a supply shock capable of causing movements in aggregate demand of even greater magnitude.

In order to provide answers to the the second question, i.e., to assess the lives versus livelihoods tradeoff, the prevailing modelling strategy consists in incorporating a SIR-type epidemiological model into a pre-existing macroeconomic framework, obtaining a macro-epidemiological model. Macro-epi models differ in terms of granularity and focus. In this context, it is generally recognized that the decentralized (competitive) equilibrium is not Pareto optimal: albeit individuals engage in endogenous social distancing to protect themselves from the risk of being infected, they do not fully internalize the *infection*

³See, for example, the huge number of columns at vox.eu and the books by Baldwin and Weder di Mauro (2020a) and Baldwin and Weder di Mauro (2020b).

externality, i.e. the effect of their socializing on the risk of infection for their contacts. In other words, endogenous or voluntary social distancing is not effective enough in terms of epidemic containment. This premise makes room for forced social distancing, i.e., public health measures such as the lockdown.

Eichenbaum et al. (2020) build a SIR-RBC model and show that containment policies are optimal when they are fine tuned to the magnitude of the infection. Acemoglu et al. (2020) adopt a more granular approach introducing age heterogeneity. In this case, optimal containment policies take the form of age-targeted quarantines, more stringent for the old, milder for the young. Optimal epidemic containment policies are studied also by Alvarez et al. (2020). Bodenstein et al. (2020) and Krueger et al. (2020) focus on the effects of Covid-19 on the supply side and in particular on the role of the industrial structure in the transmission of the shock. Bodenstein et al. (2020) distinguish between core and non-core industries: firms in the core industries produce final goods using intermediate inputs purchased at firms in the non-core industries. An uncontained epidemic has a huge impact on labour supply and a disproportionate impact on core industries. Forced social distancing not only saves lives but it also slows down the contraction of labour supply and mitigates the negative impact on core industries. Krueger et al. (2020) adopt a more granular approach considering a multi-sector economy. In their setting, households spontaneously reallocate consumption from sectors more exposed to the infection to sectors less exposed. Hence endogenous social distancing is effective and this reduces the need to resort to stringent containment policies. The role of optimal policies in economies affected by the pandemic is explored in depth in Assenza et al. (2020).

As we anticipated in the introduction, there is a growing body of literature which employs granular tools (e.g., very detailed input-output models, production networks, agent based models) to answer Covid-related research questions. For example Bellomo et al. (2020) build a very detailed multi-scale model of the interactions of different “entities” (viruses, humans, populations) that make the disease pervasive. del Rio-Chanona et al. (2020) employ detailed input-output production networks to evaluate the trickling down pattern of demand and supply shocks in the USA. The agent based model built by Poledna et al.

(2020) has been used to assess the effects of Covid-19 in Austria. Basurto et al. (2020) provide an AB macro-epidemiological model based on the Eurace@Unibi framework and calibrated on German data.

3 Environment

The model economy we construct for our analysis is populated by firms, households, the banking system and the public sector. There are N_F firms which fall in 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. Hence the cardinality of the set of C-firms is $N_F^c = N_F^b + N_F^l$. K-firms use labor to produce capital goods which are sold to C-firms. C-firms employ labour and capital to produce consumption goods which are sold to households. Firms exit the economy when their equity becomes negative. Exiting firms are replaced by entrants so that the total population of firms (and the composition of the corporate sector) is constant.

There are N_H households which fall into two categories: N_W workers and N_F firm owners. Since the number of firms is constant, also the number of firm owners is constant. In normal times also the population of workers is constant. We characterize the scenario of **Normal Times (NT)** as a setting in which people can get sick but the (non-infectious) disease is non-lethal.⁴ Therefore all the sick will eventually recover and the total population is constant.

During an epidemic, on the contrary, population may decline because the (infectious) disease is potentially lethal. For simplicity we assume that the epidemic spreads only among workers, who represent the vast majority of the population of households. During an epidemic, therefore, the number of workers will decrease because of deaths. Since, due to the time-scale of the simulations we consider, we have ruled out the *entry* of newly

⁴As outlined below, the sole purpose of the ‘normal’ disease is to generate a baseline level of demand on the healthcare system.

born agents, the population will decline.⁵

Workers can be either active or inactive: $N_W = N_A + N_I$. The population of active workers is the labour force. The labour force, in turn, is the sum of employed (N) and unemployed (U) workers: $N_A = N + U$. The inactive population consists of old (retired) workers and sick people. When a worker gets a disease (whether ‘normal’ or epidemic) she becomes inactive and remains inactive until the end of the disease.⁶

Each active worker supplies labor on the labor market. The market for labour is characterized by *search and matching*: unemployed workers search for a job and stop searching when a match occurs with a vacancy posted by a firm. When employed, the worker earns a wage. If unemployed, she will receive an unemployment subsidy. Workers (both employed and unemployed) who get sick receive sick-pay. Retired workers receive pensions. There are N_F firm owners.⁷ Firm owners earn dividends proportional to the firm’s profit if the latter is positive. Wages as well as firm and bank profits are taxed. Unemployed, sick and retired workers do not pay taxes.

Households (both workers and firm owners) are consumers, i.e., buyers on the market for consumption goods (C-goods). The market for C-goods is characterized by search and matching: households search for the goods they want at firms and stop searching when a match occurs. Also the markets for capital goods (K-goods) is characterized by search and matching: C-firms search for the machines they want at K-firms and stop searching when a match occurs.

Households can also be savers, i.e., they may spend less than their income. Savings are employed to accumulate financial wealth in the form of deposits at banks. By assumption households do not borrow.

Firms register a financing gap when outlays (to pay wages and K-goods) are greater than liquidity.⁸ By assumption the only source of external finance are bank loans. Hence firms which cannot self-finance their outlays demand bank loans.

⁵Of course also the composition of the population will change because, by assumption, firm-owners do not catch the disease.

⁶In normal times the disease ends with recovery, during an epidemic it may end either with recovery or with death.

⁷By assumption, there is one owner per firm.

⁸By assumption firms hold liquidity in the form of deposits at banks.

For simplicity, the banking sector is assimilated into a single bank. Households and firms hold deposits at the bank. The interest rate on deposits is a fraction of the risk free interest rate (i.e., the interest rate on Government bonds). The bank extends loans to firms which need to fill the financing gap. The bank sets the interest rate on loans and the quantity of credit supplied to firms. The decision of the bank is based on the assessment of the borrowing firm's financial fragility, which is a proxy of the credit risk run by the bank. 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, a firm may face a limit on the total amount of credit it can receive (*credit rationing*).

The public sector collects taxes on wage income and profits and provides transfers in the form of unemployment subsidies, sick-pay and pensions. Government expenditure consists of public provision of healthcare services. In case of a public sector deficit, the Government issues bonds and sells them to the bank. The interest rate on Government bonds is the risk free rate. Figure 1 depicts agents' interactions on six markets: deposits, credit, labor, K-goods, consumption goods (L-goods and B-goods).

In the following, undated letters will denote exogenous variables or parameters. Dated letters will denote time-varying variables. In the macroeconomic sub-model, the time unit is a month: all decisions and market interactions take place at monthly frequency. The epidemiological sub-model runs at the basic frequency of one week, with every month in the model containing exactly four weeks. In the macroeconomic submodel, therefore, the time subscript t indicates a month while the epidemiological sub-model, the time subscript τ indicates a week.

4 The macroeconomic sub-model

4.1 Households

Households (workers and firm owners) will be indexed by $h = 1, 2, \dots, N_W, N_W + 1, \dots, N_H$ where, of course, $N_H - N_W = N_F$. In words: the first N_W households are workers, the

Figure 1: Agents and markets

last N_F households are firm owners.

4.1.1 Workers

The h -th household is a worker if $h \in (1, N_W)$. Workers can be employed, unemployed, sick or retired. Each active worker supplies 1 unit of labour inelastically. If employed, she receives the uniform nominal wage w_t and pays 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 among the population of firms. Since the wage is uniform across firms and labour is homogeneous, the worker may be employed in any of the B, L and K firms. Once an unemployed worker finds a firm with an unfilled vacancy she stops searching and the match occurs. Unemployed workers who have not succeeded in finding a job (because firms in their subset did not post vacancies or because they had already filled

all the vacancies), receive an unemployment subsidy from the Government equal to a fraction of the wage: $s_u w_t$. A sick worker receives a sick-pay $s_s w_t$. Each inactive 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.

4.1.2 Firm owners

Let's index all the firms with $f = 1, 2, \dots, N_F^b, N_F^b + 1, \dots, N_F^c, N_F^c + 1, \dots, N_F$ where $N_F^c = N_F^b + N_F^l$ and $N_F = N_F^c + N_F^k$. 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. The h -th household is a firm owner if $h \in (N_W + 1, N_H)$ where $N_H = N_W + N_F$. The household indexed with $h = N_W + f$ is the owner of the f -th firm. The income of this household consists in dividends paid by the firm it owns. Dividends, 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. Whenever a firm's equity becomes negative, it is assumed to exit the economy and another one will replace it. We assume that the initial equity of the entrant firm is provided by the household that owned the exiting firm. The owner's wealth, therefore, will be reduced correspondingly.

4.1.3 Households as consumers

In order to understand households' consumption behaviour we must start from the definition of disposable income. Each household receives income and interest payments. For simplicity only the income of employed workers and firm owners is taxed. Income sources differ from one set of households to another. Let's focus first on workers. Employed workers earn the wage w_t taxed at the rate t_w . Unemployed, sick and inactive (retired) workers receive a subsidy $s_n w_t$, equal to a fraction (the replacement rate) $0 < s_n < 1$

($n = u, s, p$) of the current wage. Interest payment accruing to the household in t are computed applying the interest rate on deposits r_d to deposits accumulated until the end of the previous month: $r_d D_{t-1}$. Hence the current disposable income of the h -th worker – i.e., household $h \in (1, N_W)$ – is:

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

Finally, firm owners receive income as dividends from the firm they own and a fraction of the dividends paid out by the bank. Dividends are a fraction of after-tax profits. Hence the disposable income of the h -th firm-owner – i.e., household $h \in (N_W + 1, N_H)$ – is

$$Y_{h,t} = (1 - t_\pi)\omega \left(\pi_{h,t-1} + \frac{1}{N_F} \pi_{b,t-1} \right) + r_d D_{h,t-1} \quad (2)$$

To decide how much to consume of each C-good, the household proceeds in four steps. First, the household estimates *human wealth* $\bar{Y}_{h,t}$ using an adaptive mechanism: $\bar{Y}_{h,t} = \xi_Y \bar{Y}_{h,t-1} + (1 - \xi_Y)Y_{h,t}$ where $\xi_Y \in (0, 1)$ is a memory parameter and $Y_{h,t}$ is current disposable income. Human wealth is a weighted average of past disposable incomes with exponentially decaying weights.

Second, the household determines the desired *budget allocated to consumption*: $C_{h,t} = \bar{Y}_{h,t} + \chi D_{h,t-1}$ where $D_{h,t-1}$ is the household's non-human (i.e. financial) wealth (deposited at the bank) and $\chi \in (0, 1)$ is the fraction of wealth devoted to consumption. If the consumer does not receive income - for instance because a firm owner does not receive dividends - she will de-cumulate her financial wealth to form a consumption budget.

Third, the consumer allocates a fraction c^b of the consumption budget to the consumption of basic goods (B-goods hereafter). Hence the fraction $(1 - c^b)$ of the budget will be allocate to luxury goods (L-goods). We assume that c^b is a decreasing function of the

price of B-goods relative to L-goods. In symbols: $c^b = \frac{N_F^b P_{t-1}^l}{N_F^c P_{t-1}^b}$ where P^b (resp: P^l) is an aggregator of the individual B-prices (L-prices).⁹ There may be periods in which the household's consumption budget is larger than available liquidity (i.e., deposits plus current income). In this case, the consumer is liquidity constrained; she will first spend on basic goods and then allocate any remaining funds beyond the desired consumption of basic goods to the consumption of luxury goods.

Fourth, the consumer visits C-firms in order to purchase goods. C-firms will be indexed by $i = 1, 2, \dots, N_F^b, N_F^b + 1, \dots, N_F^c$ where, of course, $N_F^c - N_F^b = N_F^l$. In words: the first N_F^b firms are B-firms, the last N_F^l firms are L-firms. In the following we will deal with B-firms but exactly the same market protocol is applied to L-firms.

On the market for B-goods there are N_F^b firms indexed with $i \in (1, N_F^b)$. The h -th consumer visits z_c B-firms, ranks them in ascending order of posted price and demands consumption goods starting from the firm charging the lowest price. If she does not exhaust the consumption budget devoted to B-goods ($c^b C_{h,t}$) at the first firm, the consumer will move up to the second firm in the ranking and so on. This implies that there is an implicit *negative elasticity* of the demand for the good produced by the i -th B-firm to the relative price $\frac{P_{i,t}^b}{P_t^b}$ where $P_{i,t}^b$ is the price charged by the i -th B-firm and P_t^b is the average price of B-firms.

Contrary to Assenza et al. (2015) and similar previous models, in this paper we assume that the consumer keeps memory of the firms she visited in the past. We assume that in period t the set of visited firms consists of the firm the consumer visited in $t - 1$ at which she purchased the largest amount of goods, and of $z_c - 1$ firms selected randomly. In this way, current demand for a firm's output becomes a more reliable signal of future demand.

If, at the end of her visits to B-firms, the household has not spent the consumption budget allocated to B-goods, she will save involuntarily. Denoting with $C_{h,t}^b$ actual expenditure

⁹If 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. Ideally, if every C-firm sets the same price, it should be entitled to the same fraction of the consumption budget.

on B-goods, involuntary saving will be $S'_b = c^b C_{h,t} - C_{h,t}^b$. The search and matching mechanism leads to the coexistence of queues of unsatisfied consumers (involuntary savers) at some firms and involuntary inventories of unsold goods at some other firms.

The market protocol for L-goods is the same. The consumer visits a subset z_c of L-firms and starts purchasing goods from the firm which posts the lowest price. If the budget allocated to L-goods $(1 - c^b)C_{h,t}$ has not been entirely spent, she will add the residual to her savings. Denoting with $C_{h,t}^l$ actual expenditure on L-goods, involuntary saving will be $S'_l = (1 - c^b)C_{h,t} - C_{h,t}^l$.

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 on B-good and L-goods:

$$S_{h,t} = Y_{h,t} - C_{h,t} + S'_b + S'_l = Y_{h,t} - C_{h,t}^b - C_{h,t}^l \quad (3)$$

Savings are used to accumulate financial wealth in the form of deposits. The law of motion of financial wealth for the h -th household therefore is $D_{h,t} = (1 + r_d)D_{h,t-1} + S_{h,t}$. By assumption consumers do not hold Government bonds.

4.2 Firms

4.2.1 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 will describe the behaviour of a generic C-firm. The firm has some *market power* on its own local market (i.e., there are as many local markets as there are firms). It has to set individual price and quantity under uncertainty. It knows from experience that if it charges higher prices it will get 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. The best the firm can do in this setting consists in charging a price

as close as possible to the average price and producing 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 i -th C-firm, $i = 1, 2, \dots, N_F^b, N_F^b + 1, \dots, N_F^c$ must choose in t the price and desired output for $t + 1$, i.e., the pair (P_{it+1}, Y_{it+1}^*) . Desired output is determined by expected demand $Y_{it+1}^* = Y_{it+1}^e$. The firm's information set in t consists of (i) the relative price $\frac{P_{i,t}}{P_t}$ – where P_{it} is the price of the i -th good and P_t is the average price level – and (ii) excess demand

$$\Delta_{it} := Y_{it}^d - Y_{it} \quad (4)$$

where Y_{it}^d is the demand for the i -th good and Y_{it} is actual output. Δ_{it} 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.

A firm can decide either to update the current price or to vary the quantity to be produced. The decision process is based on two *rules of thumb* which govern price changes and quantity changes respectively.

The *price adjustment* rule is:

$$P_{i,t+1} = \begin{cases} P_{i,t}(1 + \mathbf{1}_u \rho_p) & \text{if } \Delta_{i,t} > 0 \\ P_{i,t}(1 - \mathbf{1}_o \rho_p) & \text{if } \Delta_{i,t} \leq 0 \end{cases} \quad (5)$$

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

Analogously $\mathbf{1}_o$ is equal to 1 if the firm has overpriced (i.e., if $\frac{P_{i,t}}{P_t} > 1$), 0 otherwise.

Excess demand Δ_{it} and the relative price $\frac{P_{i,t}}{P_t}$ 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:

$$Y_{it+1}^* = Y_{it+1}^e = \begin{cases} Y_{i,t} + \mathbf{1}_o \rho_q \Delta_{i,t} & \text{if } \Delta_{i,t} > 0 \\ Y_{i,t} + \mathbf{1}_u \rho_q \Delta_{i,t} & \text{if } \Delta_{i,t} \leq 0 \end{cases} \quad (6)$$

where ρ_q is a positive parameter, smaller than one. $\mathbf{1}_o$ and $\mathbf{1}_u$ are the indicator functions define above. Both the direction and the magnitude of quantity adjustment are determined by excess demand. The firm will increase (reduce) the quantity produced in $t + 1$ if it has overpriced (underpriced) the good and experienced excess demand (supply).

Technology is represented by a Leontief production function: $Y_{i,t} = \min(\alpha N_{i,t}, \kappa x_{i,t} K_{i,t})$ where α and κ represent labor and capital productivity respectively and $x_{i,t} \in [0, 1]$ is the rate of capacity utilization at firm i . When capital is employed at full capacity – i.e. when $x_{i,t} = 1$ – output will be $\hat{Y}_{i,t} = \kappa K_{i,t}$. This is “full capacity” output. Given a stock of undepreciated capital, actual capital in $t + 1$, K_{it+1} is given – being determined by investment carried out in t , I_{it} (to be discussed momentarily) – and cannot be modified in $t + 1$. Hence in period $t + 1$ the maximum attainable output is $\hat{Y}_{it+1}$.

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): $K_{it+1}^* = Y_{it+1}^*/\kappa$. If actual capital is greater than the capital requirement, the desired *rate of capacity utilization* will be smaller than one. If actual capital is smaller than the capital requirement, the former will be utilized at full capacity (the rate of capacity utilization will be one) but desired output will not be reached.

Whatever the scenario, if actual employment in t N_{it} is smaller than labor required to reach the feasible level of activity in $t + 1$, the firm will post vacancies. If the opposite holds true the firm will fire workers.

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:

$$w_{t+1} = \begin{cases} w_t [1 + u_{up} (u^T - u_t)] ; & u_t - u^T > 0 \\ w_t [1 + u_{down} (u^T - u_t)] & u_t - u^T < 0 \end{cases} \quad (7)$$

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.

As mentioned above, the firm determines in t the capital stock which will be available for use in production in $t + 1$ by means of investment I_{it} . 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 $\bar{K}_{it}$. 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 $t - 1$ 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 γ . Hence investment necessary “on average” to replace worn out capital is $\frac{\delta}{\gamma} \bar{K}_{it}$.

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 $K_{it+1}^T = \frac{1}{\bar{x}} \bar{K}_{it}$ where $\bar{x} \in (0, 1)$ is the desired long run capacity utilization rate. Net investment is $K_{it+1}^T - K_{it-1}$. Therefore gross investment in t is:

$$I_{it} = \left(\frac{1}{\bar{x}} + \frac{\delta}{\gamma} \right) \bar{K}_{it} - K_{it-1} \quad (8)$$

Once investment has been determined, the i -th C-firm visits a subset z_k of K-firms. As on the market for consumption goods, this subset of K-firms always includes the previously visited K-firm from which i bought the largest amount, with the rest chosen at random. Visited K-firms are ranked in ascending order of price and the C-firm starts buying capital goods from the K-firm which has posted the lowest price. If this purchase does not exhaust planned investment, the C-firm will purchase capital goods also at the second firm in the ranking and so on. If the C-firm’s demand for K-goods has not been completely satisfied after z_k visits, it is forced to “save” the unspent portion of the

investment budget. Therefore actual investment may turn out to be lower than planned investment.

4.2.2 K-firms

In setting the price of capital goods, K-firms follow the same heuristic adopted by C-firms (see equation (5)). Denoting with P_{jt} the individual price and P_t^k the average price of capital goods, we define the indicator functions. $\mathbf{1}_o$ is equal to 1 if the firm has overpriced the good in t (i.e., if $\frac{P_{jt}}{P_t^k} > 1$), 0 otherwise. Analogously, $\mathbf{1}_u$ takes value equal to 1 if the firm has underpriced the good (i.e., if $\frac{P_{jt}}{P_t^k} < 1$), 0 otherwise. The price adjustment rule therefore is

$$P_{j,t+1} = \begin{cases} P_{j,t}(1 + \mathbf{1}_u \rho_p) & \text{if } \Delta_{j,t} > 0 \\ P_{j,t}(1 - \mathbf{1}_o \rho_p) & \text{if } \Delta_{j,t} \leq 0 \end{cases} \quad (9)$$

where ρ_p is a random positive number.

The quantity adjustment rule departs from the one adopted by C-firms (see equation (6)) 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 and sold in the future. The quantity adjustment rule of the j -th K-firm, $j = 1, 2, \dots, N_F^k$ therefore is:

$$Y_{jt+1}^* = Y_{jt+1}^e - Y_{jt}^k = \begin{cases} Y_{jt} + \rho_q \mathbf{1}_o \Delta_{jt} - Y_{jt}^k & \text{if } \Delta_{jt} > 0 \\ Y_{jt} + \rho_q \mathbf{1}_u \Delta_{jt} - Y_{jt}^k & \text{if } \Delta_{jt} < 0 \end{cases} \quad (10)$$

where Y_{jt+1}^* is the desired scale of activity, Y_{jt+1}^e is expected demand, Y_{jt}^k is the fraction of the inventory of capital goods held by firm j at time t which can be used to face demand in $t + 1$, Δ_{jt} is excess demand.

Y_{jt}^k is computed applying a rate of depreciation δ^k to the stock of unsold machine tools accumulated until t . In contrast to the quantity decision-rule, K-firms' price-adjustment rule is exactly equivalent to that of C-firms. As with C-firms, K-firms are allowed to update their price-quantity decision once per month on average. Since K-firms are endowed with a linear production function whose only input is labour, once the price-quantity

configuration has been set, a K-firm may post vacancies or fire workers in order to fulfill labor requirements.

4.3 The banking system

Once the quantity to be produced has been set and the cost of inputs determined, the firm has to deal with financing. Consider a generic firm, indexed by $f = 1, 2, \dots, N_F$. If the firm's internal liquidity (i.e., the current deposits held at the bank) D_{ft-1} is “abundant”, i.e., 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

$$F_{ft} = wN_{ft} + \mathbf{1}_c P_{t-1}^k I_{ft} - D_{ft-1} \quad (11)$$

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 the firm assesses the financing gap (and the demand for loans) before accessing the market for capital goods. Hence machinery and equipment 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). As we will see momentarily, both decisions will be affected by the borrower's leverage λ_{ft} :

$$\lambda_{ft} = \frac{L_{ft}}{E_{ft} + L_{ft}} \quad (12)$$

where L_{ft} is the firm's debt and E_{ft} is equity or net worth.

The interest rate charged by the bank to each firm is determined as a mark up μ on the risk free interest rate r . Adopting the expression pioneered by Bernanke et al. (1996), the firm is charged an *external finance premium* increasing with the probability of default

which in turn is (non-linearly) increasing with leverage. The bank makes an assessment of probability of default for the i -th C-firm, which is given by:

$$p(\lambda_{it}) = \frac{e^{b_{0c}+b_{1c}\lambda_{it}}}{1 + e^{b_{0c}+b_{1c}\lambda_{it}}} \quad (13)$$

Analogously, the assessed probability of default for the j -th K-firm is:

$$p(\lambda_{jt}) = \frac{e^{b_{0k}+b_{1k}\lambda_{jt}}}{1 + e^{b_{0k}+b_{1k}\lambda_{jt}}} \quad (14)$$

In the end, therefore the interest rate charged to the generic f -th firm is a function of the risk-free interest rate and of the firm's leverage:

$$r_{ft} = \mu f(r, \lambda_{ft}) \quad (15)$$

where the function $f(\cdot)$ is increasing with all the arguments.¹⁰

In order to determine the size of the loan, the bank first sets a tolerance level for the potential loss Γ_b on credit extended (to any borrower) as a fraction ϕ of its net worth: $\Gamma_b = \phi E_{bt}$. The borrower's total debt in t will be $\Phi_{ft} + L_{ft-1}$ where Φ_{ft} 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: $(\Phi_{ft} + L_{ft-1})p(\lambda_{ft}) = \phi E_{bt}$. Therefore the new credit line is:

$$\Phi_{ft} = \frac{\phi}{p(\lambda_{ft})} E_{bt} - L_{ft-1} \quad (16)$$

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

$$\dot{L}_{ft} = \min(\Phi_{ft}; F_{ft}) \quad (17)$$

¹⁰For the specification of $f(\cdot)$ see Assenza et al. (2015).

i.e., the minimum between the new credit line and the financing gap, which is equal to the demand for loans. If the latter is greater than the former the firm will be rationed on the credit market and therefore forced to scale down production. Finally, firms in each period repay a fraction ζ of the total debt to the bank.

4.4 Net worth updating

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

$$E_{ft+1} = E_{ft} + (1 - t_\pi)(1 - \omega)\pi_{ft} \quad (18)$$

where t_π is the tax rate on profits, ω is the dividend payout ratio and π_{ft} is the firm's profit. Whenever the firm's equity turns negative, the firm exits the economy and is replaced by a new firm. The owner of the exiting firm confers the initial net worth of the entrant firm (out of her own private wealth). Hence, the population of firms is kept constant. If a firm's liquidity (its deposits) are smaller than zero at the end of the period, 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 deposits become zero, but the firm does not exit the economy unless its equity is also negative.

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

$$E_{bt+1} = E_{bt} + (1 - t_\pi)(1 - \omega)\pi_{bt} - BD_t \quad (19)$$

where π_{bt} is the bank's profit and BD_t is *bad debt*, i.e., the book value of non-performing loans. We assume that the bank remunerates deposits and earns interests on loans (if borrowers are solvent) and on Government bonds. The interest rate on deposits is determined by marking down the risk-free interest rate.

4.5 The Public Sector

The public sector raises tax revenues on wage income and profits $TA_t = t_w w_t N_t + t_\pi \left(\sum_{f=1}^{N_F} \pi_{f,t-1} + \pi_{b,t-1} \right)$. The Government extends unemployment subsidies, sick-pay and pensions to workers and makes interest payments on outstanding Government bonds to the bank. Total transfers therefore are $TR_t = (s_u U_t + s_p N_{I,t} + s_s \mathcal{I}_t) w_t + r B_{t-1}$ where $\mathcal{I}_t$ denotes the number of sick people (who receive sick-pay) in month t ¹¹ and B_{t-1} is the outstanding public debt at the end of period $t-1$. Government expenditure consists in providing of *health care services* to each citizen. This expenditure is assumed to be a constant fraction of full employment output, taking the initial population of active agents as a basis for calculation. In symbols:

$$G^h = g^h \alpha N_{A,0} \quad (20)$$

This amount is spent on the output of both K-firms and C-firms and may be taken to represent both the purchase of equipment necessary for healthcare provision and consumption of medical workers (who are not explicitly modelled here). The expenditure is in the first instance allocated to firms according to their relative revenue in the previous period: the 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 total revenue. If, after the first round of expenditure distribution, the government has been unable to spend the entire amount G^h (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^h has been spent. It is assumed that this expenditure is directly converted into a capacity of the healthcare system to provide services, i.e., a supply of healthcare equal to G^h .

A public sector deficit occurs when taxes turn out to be lower than the sum of transfers and government expenditure. In this case, the government will issue new bonds. For simplicity, we assume that the Government sells its bonds only to the bank. We assume moreover that regulation (a portfolio constraint) forces the bank to purchase government

¹¹During an epidemic only infected people who are tested and detected will become inactive and receive sick-pay.

bonds.

4.6 Demand and supply of healthcare

As mentioned in the introduction, one of the strengths of the AB approach is the capability of encompassing processes that run at different time frequency. We exploit this property to capture an important real world feature: diseases (and especially epidemics) evolve at a much higher frequency than the macroeconomy. 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 “normal” – i.e., non-infectious and non-lethal – disease with a certain probability. The presence of this disease in turn generates a demand for healthcare services. This “normal” demand will provide the baseline against which we will evaluate the effects on the healthcare system of the additional pressure generated by the epidemic.

For simplicity we assume that only workers (both active and inactive) may get sick. Since the normal disease is non-infectious, the sick cannot infect the healthy: in normal times there are no spillover effects of the disease.¹² The i -th sick agent generates a demand of healthcare in week τ – denoted with $h_{i,\tau}^i$ – (hence a pressure on the healthcare system) which is increasing with age (age_i) and affected by an idiosyncratic shock:

$$h_{i,\tau}^d = h_1 age_{i,\tau} + h_2 u_{i,\tau} \quad (21)$$

where $u_i \sim \mathcal{U}(0, 1)$. We have divided the population in three segments, the young (15% of the population), the middle-aged (65%) and the old (20%). These numbers roughly capture the current composition by age of the population of Lombardy. The variable age_i assumes values 1 if the agent is young, 2 if middle-aged and 3 if old. The total demand for healthcare is $H_\tau^d = \sum_i^{\mathcal{I}_\tau} h_{i,\tau}^d$ where $\mathcal{I}_\tau$ is the number of sick people in week τ . The

¹²The “normal disease” is not intended to depict any actually existing illness; its purpose is solely to create a baseline level of demand on the healthcare system.

healthcare services actually provided H_τ is the minimum between the demand and the supply G^h (see equation (20)):

$$H_\tau = \min(H_\tau^d, G^h) \quad (22)$$

Hence the demand for healthcare may be rationed. If an agent becomes sick and enters the healthcare system to receive treatment, its demand is added to that of those already receiving healthcare, and if that agent's demand exceeds the remaining supply, the agent is rationed and receives only a fraction of the desired supply of healthcare (or, in the extreme case, none at all).

Sick agents who were previously economically active become inactive and receive sick-pay. Retired agents who become sick will continue to receive pension payments. The normal disease is not lethal: after a fixed number of weeks (4 in the present calibration) the sick recover. In the case of the normal disease, recovery does not imply immunity: recovered agents may randomly become susceptible again in the future. This assumption implies that the normal disease will not die out.

5 The epidemiological sub-model

In this section we describe the dynamics of an epidemic, i.e., an *infectious* disease. The epidemic differs from the normal disease because of the transmission from one subject to the others through *contagion*.

5.1 The beginning: what happens to Patient Zero?

In this section we describe the initial step of the contagion with the help of figure 2.

Figure 2: The initial stage of the epidemic

At a certain point in time (say week zero), a (small) number of workers¹³ get the infectious disease. These people are the *initial infected* (and infectious) and will be denoted with $\mathcal{I}_0$.¹⁴ The (healthy and) susceptible agents after the appearance of the infected are $\mathcal{S}_0 = N_{W,0} - \mathcal{I}_0$ where $N_{W,0}$ is the initial population (of workers). In figure 2 we represent the initial infected and susceptibles as two branches of the tree originating from the initial population.

Some infected agents develop mild symptoms or do not develop symptoms at all (*non-symptomatic* for short). The cardinality of this set of agents is $\mathcal{I}_0^n$. In this case the infection can be detected only if the agent is subjected to a test. We assume that in each week the infection of non-symptomatic agents will be detected with a (small and

¹³For simplicity, firm owners are not involved in the epidemics.

¹⁴In the simulations, we set $\mathcal{I}_0 = 5$. The fraction of the infected in the initial population therefore is $5/2500 = 2/1000$. Since, as outlined below, the spread of the disease is partly stochastic, this was found to be the minimum number of initial infected necessary to generate a simulated epidemic which does not die out immediately given the present configuration of numerical parameters.

constant) probability π^r .¹⁵ We will denote the number of detected non-symptomatic agents with $\mathcal{I}_{d,0}^n$. These agents are quarantined and therefore cannot spread the disease. Non-symptomatic and un-detected infected people (whose number is $\mathcal{I}_{u,0}^n$) can still spread the disease.

People who develop serious symptoms ($\mathcal{I}_0^s$) are detected with certainty. The probability for an agent to develop serious symptoms is increasing with age.¹⁶ All agents whose infection is detected (both non-symptomatic and seriously ill) – which we denote with $\mathcal{I}_0^d = \mathcal{I}_{d,0}^n + \mathcal{I}_0^s$ – 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 require hospitalisation and therefore express demand for healthcare services according to equation (21). For simplicity, we assume that people who are quarantined at home do not need healthcare. People quarantined at home in week zero will recover (with certainty) in one of the following weeks which we denote with $0 + D_d$ where D_d is the duration of the disease (we will be more precise momentarily). Also non-symptomatic undetected agents (not quarantined) eventually recover following the same pattern. The only agents who may die from the disease are hence those who develop serious symptoms.

5.2 The contagion spreads through networks

At the end of the initial week, there are $\mathcal{I}_{u,0}^n$ people still capable of spreading the disease in week 1. In period 1 therefore there will be new infected people.

The infected remain ill for a (small) number of weeks.¹⁷ During the infectious period¹⁸, if undetected, the infected can spread the disease either in the workplace or during leisure time. 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

¹⁵In the present calibration $\pi^r = 1\%$, i.e. in every week during which an agent is infected, they are detected with probability 0.01. In fact, the detection procedure in the case of mild or no symptoms may not be effective and in any case, not all infected may be tested.

¹⁶In the simulations we assume that this probability is 1% for the young, 2.5% for the middle aged and 20% for the old.

¹⁷In the present calibration, the duration of the disease for each infected agent – denoted with D_d – is drawn from a uniform distribution whose support is the interval $4 \leq D_d \leq 6$ rounded to the closest integer.

¹⁸In our calibration, the infectious period consists of the initial three weeks of illness.

to depict the transmission of the epidemic among agents. Contagion spreads in two networks: the employment network and the network of social contacts. Workers are nodes in the *employment network*. Each worker is linked to all her co-workers in the firm she works for. If a worker is infective, she can spread the contagion to her (susceptible) co-workers.

Contagion occurs also during leisure time. To capture this process we build a *network of social contacts*. Connections in this network can be either permanent or temporary. Each worker has a set of *permanent connections* consisting of family and close friends. The total number of permanent connections is a (very small) fraction of the maximum number of possible undirected connections between workers, $\frac{N_{W,\tau}(N_{W,\tau}-1)}{2}$.¹⁹ Moreover, each agent has a set of *temporary connections* (e.g., acquaintances and people occasionally encountered outside the workplace). Temporary relationships are represented by a network of contacts which is re-set every week. For this purpose, each week a random number is drawn from a normal distribution and rounded to the closest integer to set the total number of temporary connections. Pairs of agents are then drawn at random and a temporary connection is formed between them (if previously unconnected), for that week. We assume that each infected agent meets all the agents she is connected to (both at the workplace and during leisure time) in every week. Let $\mathcal{E}_\tau$ denote the set of connections in week τ which involves exactly one infected and one susceptible agent. We will denote the cardinality of this set with $N_\tau^\mathcal{E}$. We assume that only a fraction (the contagion 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

$$\dot{\mathcal{I}}_\tau = \rho_c N_\tau^\mathcal{E} \quad (23)$$

where ρ_c is the (exogenous) transmission (or contagion) rate.²⁰ We then take a sample of size $\dot{\mathcal{I}}_\tau$ from the set $\mathcal{E}_\tau$. In sampling, we assure that the likelihood of being drawn is highest for permanent connections, second-highest for workplace connections and lowest

¹⁹In the present calibration, the number of total permanent connections is 1/750 of all possible connections. We also ensure that each agent has at least one permanent connection.

²⁰In the simulations we set this parameter to $\rho_c = 0.185$.

for random connections.²¹

Each of these potential new infected leads to an actual new infected (i) with certainty in the scenario of *Uncontained Epidemic (UE)*, (ii) with a probability that may be smaller than one in the scenario of (endogenous or spontaneous) *Social Distancing (SD)*, i.e., a setting in which people voluntarily give up social interaction, or at least are more cautious in their interactions, when they feel the risk of being infected is *too high* (we will be more precise momentarily).

Since the different types of connections have different probabilities of being drawn in the sampling process, the infection rate (i.e., the ratio of the new infected to the number of connections between one infected and one susceptible) differs across types of connections, being highest for permanent connections, second highest for workplace connections and lowest for random connections.

5.3 The progression of the epidemic

Let's consider a generic week τ . The flow of new infected in this week will be $\dot{\mathcal{I}}_\tau = \rho_c N_\tau^\mathcal{E}$. Hence the cumulative number of infections at the end of the week will be $\mathcal{I}_\tau = \mathcal{I}_{\tau-1} + \dot{\mathcal{I}}_\tau$ where $\mathcal{I}_{\tau-1}$ are cumulated infections between week zero and week $\tau - 1$. New non-symptomatic agents will be $\dot{\mathcal{I}}_\tau^n$, of which $\dot{\mathcal{I}}_{d,\tau}^n$ will be detected and quarantined. Non-symptomatic and un-detected new infected people ($\dot{\mathcal{I}}_{u,\tau}^n$) still populate the networks of employment and social contacts and can spread the disease. Non-symptomatic agents recover (with certainty) “during the disease”, i.e., in one of the weeks between $\tau + D_{min}$ and $\tau + D_{max}$ where D_{min} and D_{max} are the minimum and maximum duration of the disease.²² We will denote the number of non-symptomatic agents who got sick in previous weeks and recovered in τ with $\dot{R}C_\tau^n$.

Sick people with serious symptoms in week τ will either die during the disease or recover at the end of the disease, i.e., in $\tau + D_{max}$. In each week during the disease, the i -th agent with serious symptoms will face a probability of death which is increasing with age and

²¹In the simulations, the number corresponding to $N_\tau^\mathcal{E}$ is rounded to the closest integer. The number of potential new infected may be smaller than $\dot{\mathcal{I}}_\tau$ because the sample taken from $\mathcal{E}_\tau$ may contain multiple connections involving the same susceptible agent.

²²As mentioned above, the duration of the disease ranges between $D_{min} = 4$ and $D_{max} = 6$ weeks.

with excess demand for health care. The i -th agent who developed serious symptoms in one of the weeks before τ may die in week τ with probability

$$\pi_{i,\tau}^m = \pi^m age_i + h_3(h_{i,\tau}^d - h_{i,\tau}^s) \quad (24)$$

where π^m and h_3 are positive parameters. The demand for health of agent i , $h_{i,\tau}^d$, is defined by equation (21) while the amount of healthcare they actually receive, $h_{i,\tau}^s$, depends on whether or not the healthcare system has sufficient free capacity.²³ We will denote the flow of deceased in week τ with $\dot{DE}_\tau$. By assumption the dead will not be replaced by newly born, hence there will be no bequests. When people die, their assets are simply written off. The aggregate of dead people at the end of week τ will be $DE_\tau = DE_{\tau-1} + \dot{DE}_\tau$.

If an agent with serious symptoms has not died after the maximum duration of the disease, she will recover. We will denote the flow of agents hospitalized in previous weeks and recovered in τ with $\dot{RC}_\tau^s$. Overall, the flow of recovered agents will be $\dot{RC}_\tau = \dot{RC}_\tau^s + \dot{RC}_\tau^n$. The aggregate of recovered people at the end of week τ will be $RC_\tau = RC_{\tau-1} + \dot{RC}_\tau$.

Total population will decline because of the mortality of the disease: $N_{W,\tau} = N_{W,\tau-1} - \dot{DE}_\tau$. The set of susceptibles will decrease because of the “removed” agents and because of new infections: $\mathcal{S}_\tau = \mathcal{S}_{\tau-1} - \dot{\mathcal{I}}_\tau - \dot{RC}_\tau - \dot{DE}_\tau$. Hence $N_{W,\tau} = \mathcal{S}_\tau + \mathcal{I}_\tau + RC_\tau$.

In order to contain the epidemic, the government may implement lockdown measures which are described in the simulation experiments shown below. Moreover, we introduce the possibility for agents to endogenously engage in social distancing. This experiment is also described below. If she was previously active, the recovered agent will re-enter the labour force as an unemployed agent and begin to look for a job. In contrast to the normal times scenario, we assume that, once recovered, agents become immune to the infectious disease: they do not become susceptible again for the rest of the simulation.

²³When an agent develops serious symptoms she joins a randomised queue of agents who have an excess individual demand for healthcare, i.e., for whom $h_{i,\tau}^d > h_{i,\tau}^s$.

6 Calibration and baseline simulation: Normal Times

We set a benchmark by constructing a *baseline simulation* of the model depicting **Normal Times**, that is, a simulation in which there is a normal disease but no epidemic, hence no lockdown or any extraordinary policy measures.

In order to calibrate the model for this baseline, we draw on macroeconomic data for the Lombardy region of Italy. For this purpose, we obtain data for real GDP, consumption, gross fixed capital formation and employment for Lombardy from the website of the Istituto Nazionale di Statistica (Istat).

At the regional level, data for GDP and its components are available only at annual frequency, with our time-series ranging from 1995 to 2017. Moments and statistics calculated from these data are used to calibrate and validate the model. We apply the HP filter to the empirical time series and then calculate the standard deviations (relative to the trend component) and autocorrelations of the filtered series. Table 1 shows the empirical statistics obtained in this fashion, together with confidence intervals generated by means of bootstrapping.

Table 1: Empirical evidence for Lombardy (1995-2017): Descriptive statistics

Statistic	GDP	Consumption	Investment	Employment rate
Std. deviation	0.012047 (0.011986; 0.012109)	0.010384 (0.010347; 0.010420)	0.030953 (0.030920; 0.03099)	0.00397 (0.003963; 0.003982)
1st order autocorr.	-0.045594 (-0.050028; -0.041160)	0.320343 (0.317856; 0.322830)	0.100968 (0.09894; 0.102992)	0.280675 (0.278249; 0.283102)
2nd order autocorr.	-0.123542 (-0.128026; -0.119058)	-0.39909 (-0.40286; -0.39532)	-0.170864 (-0.173760; -0.167969)	-0.258511 (-0.26116; -0.255863)

The starting point of the calibration procedure is the set of parameters presented by Delli Gatti and Grazzini (2020) who provide a Bayesian estimation of the model discussed in Assenza et al. (2015, 2018) using quarterly data for the United States. Since the (macroeconomic component of) the present model runs at a monthly frequency, the numerical values of several parameters such as the risk-free interest rate or the rate of capital depreciation have been divided by three to reflect the interpretation of one period as one month. Several parameters are fine-tuned in order to replicate with the simulated time series the moments and statistics shown in table 1. The full set of parameters is shown in table 6 in appendix A.

We run the model 100 times with different random seeds. For each individual run we then construct annual time series from the simulated monthly time-series, apply a HP filter to the simulated data and subsequently calculate the simulated equivalents of the statistics. The means and confidence intervals of the resulting statistics are shown in table 2. The model does a fairly good job at reproducing the empirical standard deviations of GDP and consumption. Simulated investment, on the other hand, is more volatile than GDP but also significantly more volatile than empirical investment. This is due, to a large extent, to the nature of our framework. Since we use a “closed economy” model, simulated GDP only consists of private consumption and investment (along with public consumption for healthcare which, as outlined above, is constant). Lombardy, on the contrary, is open to trade with other regions and the rest of the world. Hence the actual volatility of GDP and consumption in Lombardy can be jointly reproduced only if simulated investment is more volatile than empirical investment.

Similarly, the simulated employment rate is much more volatile than its empirical counterpart. The higher volatility of the employment rate is explained by the fact that in our model, employment is tied to current production more closely than in reality.

Table 2: Simulated data: Descriptive statistics

Statistic	GDP	Consumption	Investment	Employment rate
Std. deviation	0.013514 (0.013038; 0.013991)	0.010837 (0.010382; 0.011292)	0.079079 (0.076476; 0.081682)	0.012510 (0.012048; 0.012971)
1st order autocorr.	-0.003124 (-0.041005; 0.034757)	0.328187 (0.298211; 0.358162)	-0.193938 (-0.229883; -0.157992)	-0.014004 (-0.052676; 0.024669)
2nd order autocorr.	-0.274447 (-0.306511; -0.242384)	-0.292141 (-0.324140; -0.260143)	-0.266254 (-0.301892; -0.230617)	-0.266772 (-0.299029; -0.234516)

The model performs fairly well at reproducing most of the autocorrelations we consider but the first order autocorrelations of simulated employment and of simulated investment have a negative sign, contrary to the corresponding empirical autocorrelations. In the case of employment, this is due to employment being closely tied to GDP in the model. Hence the sign of the first autocorrelation of employment in simulated data is the same as that of GDP. In the real world, employment and GDP of Lombardy are less synchronized so that the first order autocorrelations of the two variables have opposite sign.

As to investment, in the model the investment decisions of firms are partly stochastic

which makes it difficult to produce positive autocorrelations at the annual frequency in a model running at a monthly frequency. However, a side effect of the negative autocorrelation of investment is that GDP is slightly negatively autocorrelated at the first lag (in line with the empirical evidence). Figure 3 shows the autocorrelations of output, consumption, investment and the employment rate up to lag 6 while figure 4 shows the cross-correlations of output, consumption, investment and the employment rate with output. Overall, the fit is comparable to that presented by Assenza et al. (2015) who work with quarterly data.

Figure 3: Empirical and simulated autocorrelations

Figure 4: Empirical and simulated cross-correlations

7 Epidemic curves: three scenarios

Having constructed and validated the baseline in which only the normal disease exists, we move on to introducing the epidemic disease and consider first the epidemic variables. We simulate three scenarios.

The **Uncontained Epidemic scenario (UE)** involves an outbreak of the infectious disease without (i) any change in the behaviour of individual agents and (ii) any countermeasures taken by the government (e.g., prohibitions of social gatherings, closures of businesses etc.). Under these circumstances, the disease can spread freely throughout the population.

It appears somewhat implausible to assume that in the wake of a serious epidemic, agents would not change their behaviour even in the absence of measures mandated by the government. For this reason, we will consider also the (voluntary) **Social Distancing scenario (SD)** in which agents spontaneously adopt “social distancing” during the epidemic. Similarly to the approach employed by Baskozos et al. (2020), this is modelled as

a discrete choice at the individual level. First we compute a *distancing index* d_τ governed by the following law of motion:

$$d_\tau = \iota d_{\tau-1} + (1 - \iota) \mathbf{N} \times \mathbf{D}_\tau, \quad (25)$$

where $\mathbf{N}$ is a row vector containing three parameters describing the intensity of choice and $\mathbf{D}_\tau$ 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 $\mathcal{I}_\tau^d - \bar{\mathcal{I}}_{SD}$ where $\mathcal{I}_\tau^d$ is the number of currently infected *and detected* individuals and $\bar{\mathcal{I}}_{SD}$ is a fixed threshold value.²⁴ The second depicts social influence and is given by $w_d - w_{nd}$, that is the difference between the share of agents which are already socially distancing (w_d) and those who are not (w_{nd}). The third is a perceived cost of social distancing (i.e., the inconvenience or disutility the agent is subject to when engaging in social distancing), denoted by c_d . We define the probability of socially distancing π_τ^d as follows:

$$\pi_\tau^d = \frac{1}{1 + e^{-d_\tau}}. \quad (26)$$

The probability to engage in social distancing is increasing in the index d . Intuitively, an agent is more likely to distance if (i) the number of detected cases goes up (because they perceive a higher probability of becoming infected), (ii) the share of agents already distancing goes up (because of the increasing social pressure), and (iii) the perceived cost of distancing goes down.

Consider agent i . To assess whether they actually distances we draw a random number x_i from a uniform distribution $\mathcal{U}(0, 1)$. The i -th agent engages in social distancing if $\pi_\tau^d > x_i$.

Social distancing has three distinct impacts on model dynamics. Firstly, it reduces the number of temporary connections. Without social distancing, if two agents are drawn randomly to form a transitory connection, this connection is generated with probability 1. If, on the other hand, at least one of the two agents is socially distancing, the connection is instead formed with probability $\pi^c < 1$.²⁵ As a consequence social distancing makes

²⁴Therefore $\mathcal{I}_\tau^d = \mathcal{I}_\tau - \mathcal{I}_{u,\tau}^n$. In the simulations we set $\bar{\mathcal{I}}_{SD} = 1$.

²⁵In the simulations, we have set this parameter to $\pi^c = 0.5$.

the network of social interactions more sparse. Agents who socially distance however still attend their workplace (if employed) and also encounter their permanent connections.

Secondly, when an infectious agent i meets a susceptible agent j , social distancing reduces the probability that an infection will result from this meeting. As we stated above, the i -th agent socially distances if $\pi_\tau^d > x_i$. The same procedure applies to agent j , i.e., she distances if $\pi_\tau^d > x_j$. We assume that, with social distancing, the meeting between these agents generates an infection with probability $\pi_{i,j}^d = 1 - \beta \mathbf{1}_{\pi_\tau^d > x_i} - \beta \mathbf{1}_{\pi_\tau^d > x_j}$ where $\mathbf{1}$ is an indicator function which takes value 1 if the condition is fulfilled (i.e., if the agent in question is socially distancing) and 0 otherwise.

Thirdly, social distancing affects agents' demand for consumption goods. We assume that if an agent engages in social distancing, she partly shifts her demand from luxury to basic goods. The first time any agent socially distances, their demand for luxury goods is hit by a negative shock while their demand for basic goods receives a positive shock. The magnitude of these shocks declines over time during social distancing. The shocks are calibrated such that in percentage terms, the demand for luxury goods declines more strongly than that for basic goods increases.

In addition to voluntary social distancing, the epidemic can be contained by means of (government mandated) lockdowns. We characterize the **Lockdown scenario (LD)** as follows:

- A fraction of firms producing luxury goods (L-firms) are shut down completely (and their production is halted) while the rest move into “smart working” with all employees working from home. Hence the lockdown first and foremost eliminates part of the connections at the workplace, making the employment network disappear at L-firms (both shut down and still active). Firms producing basic goods and capital goods (B-firms and K-firms respectively) are not shut down.
- We also assume that smart working negatively impacts the productivity of employees, so that the maximum output that (still active) L-firms can get by switching to smart work is a fraction of the production they could obtain in normal

times.²⁶ We similarly assume that the productivity of employees at K-firms goes down during the lockdown (even though they are not shut down). This assumption allows to replicate in our model the reduced availability of durable inputs into the production of final goods which follows from the disruption of supply chains, a well know stylized facts characterizing the LD scenario.

- The lockdown limits social gatherings, eliminating connections in the network of social contacts. Hence the number of both permanent and temporary connections in the LD scenario is a fraction of the corresponding number in normal times.²⁷
- The lockdown is also associated with an increased effort to detect asymptomatic cases. The detection probability (which is exogenous in the UE and SD scenarios) becomes endogenous and time varying during the lockdown. Suppose that the lockdown starts in week τ_0 of the simulation. The detection probability in week $\tau > \tau_0$ is given by $\pi_\tau^r = \pi^r + \psi(\tau - \tau_0)$ where ψ is a positive parameter. The expression in parentheses is the interval (in weeks) which has passed since the beginning of the lockdown. We assume that the probability of detection has an upper bound: $\pi_\tau^r \leq \pi_{max}^r$. The detection probability only ceases to increase once the upper bound has been reached, even if the lockdown has ended in the meantime.
- Finally, the lockdown lowers agents' perceived cost of social distancing, making it more likely that they will choose to distance endogenously.

The effects of the lockdown on the network can be gauged by comparing figures 5 and 6 which give an example of the networks of agents (encompassing all three types of connections, i.e., workplace, random and permanent) during one period in normal times and one in lockdown. The reduction of connectivity under lockdown looks dramatic even to the naked eye.

²⁶In the simulations we assume that one third of L-firms are shut down during the lockdown. The maximum output the remaining L-firms can get by switching to smart work is 90% of the production in normal times.

²⁷In the simulations we assume that only one fourth of connections survives during the lockdown.

Figure 5: Network of agents during normal times

Figure 6: Network of agents under lockdown

The actual duration of the LD depends on the epidemiological situation. The LD is triggered automatically in period τ_0 if the number of cumulative detected cases reaches an (exogenous) ceiling $\bar{\mathcal{I}}_{max}^{LD}$. If the situation does not improve, the LD ends after a maximum duration of D_{max}^{LD} weeks. If the situation improves, i.e., the average of new detected cases over the previous 2 weeks falls below a floor $\bar{\mathcal{I}}_{min}^{LD}$, the LD will be interrupted.²⁸

²⁸In our simulations $D_{max}^{LD} = 12$, i.e., 3 months. The lockdown is activated if there are more than $\bar{\mathcal{I}}_{max}^{LD} = 5$ cumulative detected cases. The lockdown is lifted if the average of new detected cases over the previous 2 weeks falls below $\bar{\mathcal{I}}_{min}^{LD} = 1$.

Once the lockdown has ended, previously closed firms are re-opened but remain in smart working mode. Newly reopened L-firms set their initial level of production equal to the average level of demand prevailing among the L-firms which were allowed to remain open. After the lockdown, each firm which has implemented home office returns to normal operations after a stochastic number of periods. Encounters between agents slowly adjust back to their previous level, as does the perceived cost of social distancing.

In figure 7 we show the simulated epidemic curves. We run the model 100 times with different random seeds and compute the mean of the simulated data for each period (a week). The time series of the generic (simulated) epidemic variable therefore is $E_\tau = \frac{\sum_{\sigma=1}^{100} E_{\sigma,\tau}}{100}$, with $\tau = 1, 2, \dots, T_E$ where T_E is the time horizon of the simulation in weeks. The top left panel shows the number of cumulative infections $\mathcal{I}_\tau$ for $\tau = 1, 2, \dots, 200$ (i.e., almost 4 years) while the top right panel shows the flow of new infections (per week), i.e. the first difference of the cumulative infections, which we have denoted above with $\dot{\mathcal{I}}_\tau$. The bottom panels show the same curves for detected cases, $\mathcal{I}_\tau^d$ and $\dot{\mathcal{I}}_\tau^d$. The shaded area around each line represents the 95% confidence interval.

Figure 7: Comparing the epidemic scenarios (weekly)

Figure 8: Comparing the epidemic scenarios (weekly)

In the uncontained epidemic (UE) scenario, cumulative infections grow exponentially for around 8 weeks and reach a plateau after around 8 additional weeks. Overall, in 4 months, 90% of the population will be infected (with the Lombardy numbers in mind, this means that if there were no lockdown, at the end of June roughly 9 million people would have been infected). The flow of new infections reaches a peak (at week 8, i.e. after two months from the outbreak of the epidemic) of close to 20% of the population. In the social distancing (SD) scenario, the initial phase of exponential growth is similar to that of UE, but (i) it takes much longer for the number of cumulative infections to reach a plateau (approximately three years) and (ii) the fraction of the population which would be infected at the plateau is much lower (50%). The flow of new infections reaches a peak earlier than in the UE case but (iii) the maximum weekly increase of infections is much lower (around 4%) and (iv) the flow converges to zero with *damped oscillations*. The oscillatory pattern of convergence – which is absent in the UE scenario – is due to the law of motion of the distancing index. As soon as infections become smaller than the threshold $\bar{\mathcal{I}}_{SD}$, people relax and no longer socially distance, but this relaxation boosts the number of infections (here comes the second wave) which overcomes the threshold again and induces people to re-enter the SD mood and so on. This (damped) cycle of relaxation and tightening of SD is the root of the waves in this scenario. These waves, however, are smaller and smaller in amplitude. The dampening factor is simply the fact that the population of susceptibles is shrinking over time.

The lockdown scenario (LD) “looks like” the uncontained epidemic case, but on a much

smaller scale. The number of cumulative infections grows exponentially for around 4 weeks and reaches a plateau after about 8 additional weeks. Overall, in 3 months, less than 10% of the population will be infected (against 90% in the uncontained scenario). If the model were a reliable description of the Lombardy case, from this scenario we could infer that at the end of May approximately 1 million people would have been infected. Notice that this number is one order of magnitude bigger than the official number for Lombardy (100 thousand) in that period. It is well known, however, that during the first wave of the pandemic only a (very small) fraction of the infected were detected. In fact, in the bottom left panel, the cumulative *detected* cases generated by the model reach a plateau of 1% of the population, in line with the official numbers. This is not a surprise since the model has been calibrated to replicate the Lombardy numbers. What the model predicts – hence its usefulness – is that the undetected cases were around 10 times the detected cases. This is in line with the rough estimate (unofficially coming from epidemiologists) circulating in the press according to which for each detected (symptomatic) case there were around 10 undetected (and mostly asymptomatic) cases.

In the LD scenario, the flow of new infections reaches a peak of 2% of the population for the total cases (top right panel) and 0.1% for the detected cases (bottom right panel), even earlier than with SD. With the lockdown, convergence to the plateau is accelerated (compared to the uncontained case). Moreover the flow of total cases converges to zero monotonically. A mild oscillatory behavior is still emerging for detected cases due to the increasing probability of infections being detected. The dynamics of fatalities emerging from simulations, shown in figure 8 are qualitatively similar to those of actual and detected infections though of course the numbers are much lower.

Figure 9 compares the empirical epidemic curves of (cumulative and new) detected infections in Lombardy (first wave) to the simulated epidemic curves from the LD scenario, where the simulated numbers have been scaled to make the absolute numbers comparable. Overall, the model does a good job at reproducing both cumulative and newly detected cases although it (i) slightly overestimates the peak of newly detected infections and (ii) underestimates the number of cases at the end of the simulation. The reason for

this is the small size of the population in the model relative to the size of the population which it supposedly represents. With only 2500 agents in the model, each artificial agent represents 4000 of the 10 million residents of Lombardy. One infected agent in the model, therefore, corresponds to 4000 infected people in Lombardy. Given this ratio, the model has considerable trouble in jointly reproducing the empirical numbers of the first wave and the subsequent pattern of steady (if low) infections for an extended period, which was the pattern observed empirically after the first wave. In the simulations, the epidemic disease almost always dies out following a lockdown strict enough to produce empirically plausible numbers for the first wave.²⁹ From a computational perspective, however, increasing the number of households beyond the current 2500 would be extremely burdensome and at a certain point a further increase would become practically infeasible. Basurto et al. (2020) address this problem by considering a model which is more strongly simplified in both the economic and epidemiological dimension.

Figure 9: Comparing empirical and simulated infection data (weekly)

Table 3 compares the actual number of (cumulative) deaths in Lombardy at the end of the first wave (third week of June) – when the flow of new deaths went down almost to zero and the aggregate stabilized – with the number of deaths (DE) at the plateau in each of the scenarios generated by the simulations (we also provide the endpoints of

²⁹A second wave in this scenario could only spring from a newly injected initial population of infected individuals and/or from recovered agents becoming susceptible again.

the confidence interval). The lockdown scenario comes very close to reproducing the empirical number and very strongly reduces the number of fatalities compared to the UE and SD scenarios. Based on the model simulations, from the UE scenario we infer that in the absence of mitigation factors (lockdown or social distancing) at the end of the first wave approximately 370 thousand people would have died in Lombardy, a number more than 20 times bigger than actual deaths. This is in particular due to the reduced strain on the healthcare system under lockdown, which is illustrated in figure 10.

Table 3: Empirical and simulated cumulated fatalities at the end of the first wave

Empirical	16570	-	-
Simulated	<i>Mean</i>	<i>Lower</i>	<i>Upper</i>
Uncontained E.	368560	360947.21	376172.79
S. Distancing	129960	121374.23	138545.77
Lockdown	18360	16523.26	20196.74

Figure 10: Demand for healthcare (weekly)

The bold flat line is the supply of healthcare as determined in section 4.6. The uncontained epidemic generates a huge and persistent excess demand for healthcare. At the peak, demand is almost 5 times the supply and excess supply lasts for 15 weeks, i.e. almost 4 months (the entire duration of the first wave). SD and the LD dramatically reduce the need for healthcare. With LD demand peaks at a much lower level and excess demand lasts for 10 weeks. Under SD the demand for healthcare also peaks at a much lower level but since SD increases the overall duration of the epidemic, the period of excess

demand lasts almost as long as in the UE scenario. Moreover, the oscillatory dynamic is also present in the demand for healthcare under SD, meaning that the healthcare system also reaches full capacity during the second wave.

8 The macroeconomic effects of the epidemic

The discussions above focused purely on the epidemiological effects of the disease. In this section we explore the macroeconomic effects of the epidemic under the scenarios presented above.

In each panel of figure 11 we represent the departures of the most important macroeconomic variables from the baseline (i.e. Normal Times) under the scenarios of uncontained epidemic (UE), social distancing (SD) and lockdown (LD). The time unit is a month. For each macroeconomic variable (with the exception of the default rate), in the corresponding panel we report the time series of the *relative change* $m_t := \frac{M_t^S}{M_t^N} - 1$, $S = UE, SD, LD$, $t = 1, 2, \dots, 120$ where M_t^S is the value of the variable M in scenario S in month t . Similarly M_t^N is the value of the same variable in Normal times in the same period. The only exception to this way of representing the departure from the baseline is the default rate. We define the default rate (DR) as the ratio of defaulting loans to average GDP *during the pre-epidemic periods*. In the corresponding panel we report the series of the *absolute change* $dr_t := DR_t^S - DR_t^N$ since there are some individual periods in some runs in which $DR_t^N = 0$ (i.e. no defaults). Since the resulting series is quite noisy, we applied a 4-period moving average filter.

Each value of the generic macroeconomic variable (and of the default rate) is the mean of 100 simulated data points generated by different random seeds: $M_t = \frac{\sum_{\sigma=1}^{100} M_{\sigma,t}}{100}$, with $t = 0, 1, 2, \dots, T_M$. The time horizon of the simulations goes from the outbreak of the epidemic (month zero) to $T_M = 120$ months but we will focus essentially on the first 24 months (the “short run”). Beyond the short run, the departure generally stabilizes around a “long run” mean.

The patterns of GDP, employment, consumption and investment are qualitatively very

similar within each scenario. In the UE scenario, GDP declines “slowly” due to the mere effects of the epidemic. Active people get sick and become inactive, employment and production shrink and firms cannot supply the same amount of goods as before. Therefore, *even in the absence of containment measures*, the epidemic inflicts a sizable *negative supply shock* to the macroeconomy impairing the production capability of firms. On top of the supply shock, the macroeconomy experiences a *negative demand shock*, which comes from two sources.

First, people who get sick or are laid off experience a significant contraction of income as they receive sick-pay and unemployment subsidies instead of wages and consume out of this reduced income.

Second, in the simulation, the consumption expenditure coming from the people who die from the disease is permanently “removed” from aggregate demand as the deceased are not replaced by newly born by assumption.³⁰ Hence production declines further.

The negative demand shock is at the root of the drop of consumption (relative to the baseline) at the beginning of the simulation horizon, sharper than the drop of GDP. Over an extended period consumption falls again but gently. In the long run GDP and consumption stabilize at a level approximately equal to 95% of the corresponding value in Normal Times, i.e., around 5% below the baseline.

In our model, by construction, investment is closely associated to production plans. Hence it is highly correlated with consumption and GDP. As expected it falls below the baseline and oscillates irregularly but shows a modest tendency to go up towards the end of the simulation horizon.

The aggregate leverage ratio increases above the baseline but shows a tendency to decline towards the end of the simulation horizon. Loan defaults as a share of GDP appear largely unaffected.

The price level declines and shows ample and persistent fluctuations below the baseline. The relative price of luxury goods in terms of consumption goods, however does not change with respect to Normal times.

³⁰The majority of deceased people consists of inactive agents who were receiving pension payments.

As to public finances, while transfers to households for income support increase, transfers for pensions decline. With unchanged tax rates, also tax revenue declines. The decline of pension outlays, however, more than offsets the increase of sick-pay and unemployment benefits and the reduction of tax revenues. Hence, government debt declines. In the long run it will stabilize around 5% below the baseline. However, since GDP also declined, the debt over GDP ratio will not depart significantly from the corresponding value in normal times.

Social Distancing magnifies the short-term negative effect on consumption and GDP. Under this scenario, in fact, agents switch from luxury to basic goods and the demand for luxury goods declines more rapidly than that for basic goods increases. This is reflected in the huge early drop of consumption. With SD, consumption decreases by 15% (with respect to Normal Times) at the very beginning of the simulation horizon (first 6 months), a phenomenon which does not occur under UE. This drop also translates into early declines in output and investment more pronounced than in the UE scenario. Moreover, the relative price of luxury goods declines strongly – due to the shock to the composition of demand for consumption goods – leading to a decrease in the overall price level. Loan defaults increase strongly as the sudden shift in consumption induced by social distancing causes many firms to over-produce and subsequently go bankrupt.

As the share of agents who die is much smaller under SD, however, consumption and GDP rebound in the second semester of the horizon and eventually regain their pre-epidemic level. Also the relative price of luxury goods bounces back. Investment eventually overcomes the Normal Times level for a while.

With SD the sizable decline of pension outlays which occurs in UE does not materialize for the obvious reason that people do not die at the same rate. However transfers to households for income support increase and tax revenues decline so that government debt increases “in the short run” (around two years). Afterwards, debt decreases because GDP bounces back leading to a decrease of transfers and an increase of tax revenues. Over the long run the debt/GDP ratio will stabilize slightly above the value it reaches in Normal Times.

The LD scenario brings about stark and dramatic macroeconomic consequences. We assume that during the lockdown, one third of L-firms are shut down and cease to produce any output. The supply shock is therefore much bigger than in UE and SD. In the absence of policy interventions, the firms which are closed down lay off their entire workforce. The unemployed will experience a severe reduction of income, which is now limited to the unemployment subsidy. Together with (voluntary) social distancing, the massive increase of unemployment (as well as the drastically reduced supply of L-goods) makes a dent in consumption, which shrinks by 15% relative to the baseline in the first quarter of the simulation horizon. Since C-firms downsize, also investment and the production of K-goods goes down. Investment drops around 30% in the same time period. Hence GDP shrinks by 20% due to the intertwined supply and demand shocks. There is also a large spike in loan defaults. These dramatic effects, however, are reabsorbed over the medium run.

During the LD, due to the drop in supply, the price of luxury goods initially jumps up, driving the aggregate price level up 2% above the baseline. After the lockdown, once all K-firms are again allowed to produce, the relative price of luxury goods declines as demand for them is still relatively low due to social distancing.

After the initial contraction, GDP bounces back. It then experiences another, smaller decline as newly reopened firms adjust their production to the new environment. The pattern of recession and recovery in the first two years after the outbreak of the epidemic hence looks somewhat W-shaped. What's more important, however, is the fact that it takes two years for GDP to return to normal. The post-lockdown recovery therefore is immediate but partial: catching up with the baseline is a lengthy process. Over a longer horizon, it is worth noting that GDP overcomes the baseline for a long period (a few years), which is due to a combination of the adaptive process of firms in forming demand expectations and a 'backlog' in consumption demand which built up during the early stage of the lockdown. A similar pattern of W-shaped recovery and lengthy catching up can be detected also for employment, consumption and investment. After the end of the LD the price level declines and stays below the baseline for some time. Government

debt keeps rising until GDP has recovered and then decreases due to the overshooting of output.

Figure 11: Economic impact of the disease under different scenarios

Figure 12 gives a sense of the short-run tradeoff between the infection rate and the magnitude of the recession implied by the different scenarios described above. For each scenario we report in the same picture the share of new infections per month in total population and the relative departure of GDP from the baseline in each scenario. The

scenario of UE is characterized by the maximum impact on infections and the minimum short-term impact on GDP. Social distancing implies a drastic reduction in the initial severity of the epidemic relative to the UE scenario at a small initial economic cost, but as described above, it also strongly increases the length of the epidemic such that at the end of the day, the reduction in terms of total infections (and also fatalities) is not as large as it appears in the initial stage. The LD scenario, on the other hand, is successful at quickly and decisively getting the epidemic under control but implies an enormous economic loss, at least in the short run. This gives an idea of the basic tradeoffs involved in reacting to the epidemic, and by varying parameters relating to the ‘intensity’ of social distancing and the severity of the lockdown, various intermediate scenarios could be constructed.

Figure 12: Epidemiological and economic impacts under different scenarios

9 The effects of stabilization policies

In this section we will consider the effects of stabilization policies aimed at mitigating the deleterious effects of the epidemic and of the lockdown on the macroeconomy. As we have seen above, even if the epidemic were not addressed by means of a lockdown, it would generate a recession. The lockdown, however, dramatically exacerbates the magnitude of the contraction in aggregate activity. Hence it makes stabilization policies absolutely necessary. Moreover most of the countries hit by the epidemic have adopted containment measures of the lockdown type so that this scenario is the closest to reality. In this section, therefore, we will compare the LD scenario with a number of scenarios characterized by different macroeconomic stabilization measures. Due to the nature of our model, we focus on short-term policy measures ultimately aimed at boosting demand and employment rather than on longer-term ones addressing the structural changes which may be necessary to face the supply disruptions induced by the epidemic and the lockdown. The range of measures we consider qualitatively replicates the provisions actually undertaken by many Governments. Since we calibrate the model on Lombardy, we will design each scenario having in mind in particular the measures taken by the Italian government. We will consider the following policies:

Layoff ban and redundancy fund (RF): The government (i) prohibits firms from firing those workers they do not need due to the LD and (ii) takes over the wage payments to these workers.³¹ Hence workers retain their jobs and the government pays their wages. In the RF scenario we assume that the policy has a duration of 12 months and that workers are paid their full wage under the programme.

Liquidity support (LS): The government subsidizes firms in order to improve their liquidity positions. We consider two scenarios. In the first one – which we will refer to as LS(1) – *all* firms receive a *one-off* transfer of equal size. In the second one – denoted with LS(2) – only *firms which experience a liquidity shortfall*³² after having visited the

³¹Policies of this type go under the name of Cassa Integrazione Guadagni (CIG) in Italy and Kurzarbeit in Germany.

³²We define the liquidity shortfall as the difference between the sum of wage payments and planned investment on the one hand, and available liquidity on the other hand.

credit market receive such a transfer, equal to the size of the liquidity gap. In this second case the policy lasts 12 months.

Credit guarantee (CG): Instead of injecting liquidity through transfers, the government provides guarantees on bank loans beyond those which firms are able to obtain on the regular credit market. With a government guarantee, the bank has an incentive to fill the firm’s liquidity shortfall at the risk-free interest rate. If the firm defaults on this loan, the loss is taken by the government rather than the bank. This policy remains in place for 12 months.

Equity injections (EI): The government provides support by injecting equity into firms with negative equity. In this case the government becomes a part-owner of that firm and receives a share of its dividend payments. We assume that the this policy lasts for a period of 12 months.

Income support (IS): The government supports workers directly by making transfer payments. We assume that for 6 months, the government makes a monthly payment of half the unemployment benefit to each worker (in addition to all other transfers already existing in the model)

Each of the following 6 figures illustrates one of the policy scenarios listed above. Each figure consists of four panels. In each panel we show the time series of the relative departure from the baseline of a macroeconomic variable (i) under LD in the absence of policy (dashed line) and (ii) under LD *cum* policy (solid line).

Let’s start from RF, shown in figure 13. When the Government adopts RF, the most striking difference with respect to the LD scenario concerns the employment rate, which does not decrease relative to the baseline (as it did in LD) and – for the duration of the policy – even increases. This is due to the fact that the policy (as in the real world) forbids firms from firing workers in the face of a drop in demand. Since changes in production are the only reason for firing or hiring workers in our simplified setting, in the RF scenario all redundancies are actually banned for a certain period. This in turn implies that employment is not only higher than in LD without policy but also higher than in the baseline. As to GDP, with RF the recession generated by the epidemic and

the lockdown is slightly less dramatic and the recovery more pronounced. This is due to the fact that workers retain all the purchasing power of an active worker even if they stay home during the lockdown. In the LD scenario without policy, unemployed workers receive only a subsidy equal to a fraction of the real wage, whereas under RF workers keep their job and receive their full wage. Hence consumption falls less than in the absence of intervention. RF leads also to a slightly stronger initial increase in government debt relative to LD. Since the policy does not impact directly on the net worth of firms in our model (to a firm in our model it makes no difference from a balance sheet perspective whether it fires a worker it does not need or if it keeps the worker and the government pays their wage), the policy has no significant impact on loan defaults

Figure 13: Impact of the Redundancy Fund (RF)

The impact of liquidity support (LS) differs strongly depending on the type of subsidy. Support in the form of a one-off injection to *all* firms (LS(1), figure 14) appears to have a sizable positive effect while continuous support targeted to firms with liquidity shortfalls (LS(2), figure 15) is essentially ineffective. In the latter case, in fact, the time series

of each variable with and without policy overlap. In LS(2), liquidity support flows to firms which are already fragile, and only in amounts just sufficient to cover the current liquidity gaps. These firms tend to be relatively small, so that liquidity channelled to them does little to boost aggregate output. In addition, the small liquidity injections under this policy appear to be ineffective in preventing defaults. On the other hand, the small size of these transfers does not impact greatly on government debt. The one-off generalized support policy LS(1), on the other hand, has a strong effect on GDP, significantly speeding up the recovery. The generalised liquidity support improves the overall financial robustness of firms and, since it directly increases the net worth of firms, prevents a large share of defaults which would otherwise occur. Moreover, it gives firms liquid resources to service their debt and make dividend payments which in turns boosts household income and consumption.

Figure 14: Impact of one-off liquidity support (LS(1))

Figure 15: Impact of continuous liquidity support (LS(2))

The impact of credit guarantees (CG) on the pattern of GDP, shown in figure 16, is very similar to that of continuous liquidity support (LS(2)), with a marginal effect on the speed and amplitude of the recovery. Unsurprisingly, compared to LS(2), the increase in government debt is smaller since support takes the form of credit rather than direct transfers: the Government takes a loss only if a firm receiving a loan guaranteed by the government goes into default. Loan defaults increase somewhat relative to the case of no policy intervention as some firms end up defaulting on the loans guaranteed by the government.

Figure 16: Impact of credit guarantees (CG)

At first sight, the macroeconomic effects of equity support policies (EI) are surprising and somehow puzzling. As shown in figure 17, the macroeconomic outlook with equity injections appears worse than in the LD scenario, at least in the short run (around 2 years). While, as shown by the strong decrease in loan defaults, EI does prevent firm bankruptcies (at the cost of an increase in government debt), output losses are larger than in LD. The root cause of this phenomenon is, once again, the prevailing size of the firms targeted by the policy coupled with the entry mechanism embedded in the model. By assumption, the firms which are “rescued” by the government would otherwise exit the economy due to negative equity, and therefore tend to be small. In other words, EI “artificially” keeps low-output firms alive. In the present model, an exiting firm is replaced immediately by a new one whose desired output is equal to the mean output currently prevailing in the respective segment of the corporate sector. Replacement firms hence tend to be larger in terms of output than those which exit the economy. The puzzle therefore is an artefact of the firm replacement mechanism incorporated in the

model which, particularly when the model is simulated at monthly frequency, may be viewed as somewhat unrealistic. Nevertheless the result emphasises an important point, namely that rescuing failing firms alone may do little to promote recovery if those firms are not given an incentive to expand their output.

Figure 17: Impact of equity injection (EI)

Figure 18 shows that the income support (IS) scenario, whereby the government makes transfer payments to all worker households for a limited number of periods. This policy is highly effective at promoting a faster recovery from the epidemic and lockdown-induced downturn. In addition to a much smaller decline in GDP and employment, the policy counteracts the impact of the lockdown on loan defaults. It also amplifies the post-recovery overshooting effect which is also present in the scenario without policy. The faster recovery does however come at the price of a large increase in government debt necessary to finance the additional transfer payments.

Figure 18: Impact of income support (IS)

A similar degree of improvement in macroeconomic performance can be obtained by combining the one-off liquidity injection to firms (LS(1)) with the redundancy fund (RF), shown in figure 19. This policy, too, induces a large increase in government debt but appears more successful at preventing loan defaults and also strongly reduces the loss of employment particularly in the early phase of the epidemic period.

Figure 19: Impact of combined redundancy fund and one-off liquidity support

Figure 20: Cumulative output loss relative to baseline under different policies

Figure 21: Time to recovery (months) under different policies

Figures 20 and 21 present a direct comparison of the policy measures considered above, in terms of cumulative output loss relative to the baseline and time to complete recovery (in months). Tables 4 and 5 present the same information in tabular form. It can be seen that the most effective policy measures appear to be income support as well as the combination of one-off liquidity support and the redundancy fund. Both of these significantly speed up recovery and strongly reduce the output loss relative to the baseline of no policy intervention. Equity support is clearly an inferior policy option in our model, leading to additional output losses and significantly slowing down the recovery.

Table 4: Cumulative output loss relative to baseline

Policy	Median	Lower	Upper
None	-1.269208	-1.372594	-1.165822
Liquidity (1)	-0.965633	-1.032171	-0.899094
Liquidity (2)	-1.102897	-1.178526	-1.027268
Credit	-1.103067	-1.175343	-1.030791
Equity	-1.390574	-1.525336	-1.255812
RF	-1.032810	-1.097128	-0.968492
Income	-0.752603	-0.801588	-0.703618
RF + Liqu.	-0.850535	-0.904418	-0.796652

Table 5: Time to recovery (months)

Policy	Median	Lower	Upper
None	21	19.07490	22.92510
Liquidity (1)	16	15.04722	16.95278
Liquidity (2)	19	17.73600	20.26400
Credit	18	16.64337	19.35663
Equity	26	23.89264	28.10736
RF	18	16.26200	19.73800
Income	13	11.89400	14.10600
RF + Liqu. (1)	15	14.21000	15.79000

10 Conclusion

In this paper we make three contributions.

First, we contribute to the macroeconomic-epidemiological literature by proposing a new medium-sized Agent-Based macroeconomic-epidemiological model of Covid-19, which we label ABC. The model consists of a macroeconomic and an epidemiological sub-model. In both sub-models we exploit the granularity and flexibility of the agent-based approach. We put the model to work to address key research challenges brought to the fore by the widespread macroeconomic loss generated by the epidemic.

The second contribution of the paper consists in evaluating the effect of the lockdown on the trade off between saving lives and avoiding contractions of GDP. From the simulations we infer that, if uncontained, the epidemic not only has a huge death toll but it also carries a significant output loss, both in the “short run” and over an extended time horizon. In fact, after an initial mild drop GDP stabilizes around a quasi-steady state permanently below the level of normal aggregate activity. The main takeaway message of this analysis is that just letting the epidemic run to avoid sudden interruptions of economic activity does not benefit the macroeconomy (as it does not avoid a contraction of aggregate output) especially over the long run (as the economy stabilizes in a lower quasi-steady state).

Voluntary social distancing allows to reduce the death toll to a large extent but it brings about an oscillatory pattern of the epidemic curves. After the first wave, as the number of infections becomes smaller than a given threshold, people relax and no longer engage in social distancing, but this relaxation makes the number of infections grow again and

leads to a second wave, which in turn makes people more cautious and re-enter the social distancing mood. The macroeconomic loss in this case is bigger than in the uncontained epidemic case in the short run, i.e., immediately after the outbreak of the epidemic, but becomes smaller over a longer horizon.

The lockdown has a remarkable mitigating effect on the lethality of the disease. In the short run the macroeconomic loss is dramatic but GDP bounces back over a longer horizon. The trade off between lives and livelihoods is sizable in the early stage of the lockdown but fades away in the long run. Having calibrated the model on the empirical reality of Lombardy, ABC under lockdown is able to closely reproduce the number of deaths for Lombardy during the first wave. Using these simulation, we could infer that absent the lockdown, the number of fatalities in Lombardy could have been 20 times bigger than the actual figure.

The third contribution of the paper consists in carrying out a number of policy experiments to assess the efficacy of macroeconomic stabilization policies in counteracting the lockdown-induced downturn. We consider measures ranging from income and employment support to liquidity provision and credit guarantees and assess their effect on the dynamics of macroeconomic variables under the lockdown scenario. In the simulations, the most effective policy measures appear to be transfer payments to households, as well as a policy package consisting of employment support (in the form of a redundancy fund) and temporary liquidity support for firms. These measures significantly speed up recovery and reduce the output loss relative to the lockdown scenario.

References

- ACEMOGLU, D., V. CHERNOZHUKOV, I. WERNING, AND M. WHINSTON (2020): “A multi-risk SIR model with optimally targeted Lockdown,” *NBER WP*, 27102.
- ALVAREZ, F., D. ARGENTE, AND F. LIPPI (2020): “A simple planning problem for Covid-19,” *NBER WP*, 26981.
- ASSENZA, T., F. COLLARD, M. DUPAIGNE, P. FEVE, C. HELLWIG, S. KANKANAMGE, AND N. WERQUIN (2020): “The hammer and the dance: Equilibrium and optimal policy during a pandemic crisis,” *TSE Working Paper*, 20-1099.
- ASSENZA, T., P. COLZANI, D. DELLI GATTI, AND J. GRAZZINI (2018): “Does fiscal policy matter? Tax, transfer, and spend in a macro ABM with capital and credit,” *Industrial and Corporate Change*, 27, 1069–1090, <https://doi.org/10.1093/icc/dty017>.
- 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>.
- BALDWIN, R. AND B. WEDER DI MAURO, eds. (2020a): *Economics in the Time of COVID-19*, London: CEPR Press.
- (2020b): *Mitigating the COVID Economic Crisis: Act Fast and Do Whatever It Takes*, London: CEPR Press.
- 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 (2020): “Economic and epidemic implications of virus containment policies: insights from agent-based simulations,” *GROWINPRO Working Paper*, june.
- BELLOMO, N., R. BINGHAM, M. CHAPLAIN, G. DOSI, G. FORNI, D. KNOPOFF, J. LOWENGRUB, R. TWAROCK, AND M. VIRGILLITO (2020): “A multiscale model of virus pandemic: Heterogeneous interactive entities in a globally connected world,” *Mathematical Models and Methods in Applied Sciences*, 30, 1591–1651.
- 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>.
- BODENSTEIN, M., G. CORSETTI, AND L. GUERRIERI (2020): “Social distancing and supply disruptions in a pandemic,” *Federal Reserve Finance and Economics Discussion Series*, 2020-031.
- DEL RIO-CHANONA, R. M., P. MEALY, A. PICHLER, F. LAFOND, AND D. FARMER (2020): “Supply and demand shocks in the Covid-19 pandemic: An industry and occupation perspective,” *Oxford Review of Economic Policy*, 36, S94–S137.

- DELLI GATTI, D. AND J. GRAZZINI (2020): “Rising to the Challenge: Bayesian Estimation and Forecasting Techniques for Macroeconomic Agent-Based Models,” *Journal of Economic Behavior & Organization*, 178, 875–902.
- EICHENBAUM, M., S. REBELO, AND M. TRABANDT (2020): “The macroeconomics of epidemics,” *NBER WP*, 26882.
- GOURINCHAS, P. (2020): “Flattening the pandemic and recession curves,” in *Mitigating the COVID Economic Crisis: Act Fast and Do Whatever It Takes*, ed. by R. Baldwin and B. Weder di Mauro, London: CEPR Press, 31–40.
- GUERRIERI, V., G. LORENZONI, L. STRAUB, AND I. WERNING (2020): “Macroeconomic implications of covid-19: Can negative supply shocks cause demand shortages?” *NBER WP*, 26918.
- KRUEGER, D., H. UHLIG, AND T. XIE (2020): “Macroeconomic dynamics and reallocation in an epidemic,” *NBER WP*, 27047.
- POLEDNA, S., M. MIESS, AND C. HOMMES (2020): “Economic forecasting with an agent based model,” *IIASA WP*, 01.

Appendix A Parameter values

Tables 6 and 7 below provide the lists of model parameters pertaining to the macroeconomic sub-model and the epidemiological sub-model respectively.

Table 6: Macroeconomic sub-model parameters

Symbol	Description	Value
$N_{W,0}$	Initial number of workers	2500
N_F^b	Number of B-firms	100
N_F^l	Number of L-firms	150
N_F^k	Number of K-firms	50
z_c	Number of C-firms visited by consumers ³³	3
z_e	Number of Firms visited by unemployed	5
z_k	Number of K-firms visited by C-firms ³⁴	3
ξ_Y	Memory parameter for baseline human wealth	0.765
χ	Propensity to consume out of wealth	0.00825
ρ_q	Quantity adjustment parameter	0.2
ρ_p	Price adjustment random parameter	0.07
μ	Bank's gross mark-up	1.007
δ	Capital depreciation rate	0.01
γ	Probability to invest	0.15
ϕ	Bank's leverage parameter	0.0025
ζ	Debt repayment rate	0.01
ξ_K	Memory parameter for capacity utilisation	0.2
α	Labour productivity	$\frac{2}{9}$
κ	Capital productivity	$\frac{1}{9}$
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)	-15
b_{1c}	Bank's risk evaluation parameter (C-firms)	13
b_{0k}	Bank's risk evaluation parameter (K-firms)	-5
b_{1k}	Bank's risk evaluation parameter (K-firms)	5
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^h	Ratio of healthcare expenditure to full employment GDP	0.04

³³Including largest firm visited previously

³⁴Including largest firm visited previously

Table 7: Epidemiological model parameters

Symbol	Description	Value
	Share of young agents in the population	0.15
	Share of middle-aged agents	0.65
	Share of old agents	0.2
	Probability of catching the normal disease	0.0012
	Duration normal disease	4
	Susceptibility probability normal disease	0.1
	Shock to demand for luxury goods	$\frac{2}{3}$
	Shock to demand for basic goods	1.2
	Probability of serious symptoms (young, middle, old)	0.01, 0.025, 0.2
	Total number of possible connections	3123750
	Number of permanent connections	4165
	Mean of random connections per period	1115.625
	Standard deviation of random connections	124.95
	Share of deactivated L-firms in lockdown	$\frac{1}{3}$
	Constraint on production in lockdown	0.9
	Lockdown minimum duration (months)	3
	Lockdown activation threshold	5
	Lockdown lifting threshold	1
	Duration of epidemic disease (weeks)	$\ \mathcal{U}(4, 6)\ $
	Infectious period (weeks)	3
	Post-lockdown adjustment parameter	$\frac{1}{3}$
	Share of connections under lockdown	0.25
c_d	Cost of distancing	2
	Cost of distancing (lockdown)	-2
ι	Persistence of distancing index	0.7
β	Distancing effect on infection probability	0.4
ρ_c	Transmission rate	0.185
$\bar{\mathcal{I}}_{SD}$	Distancing threshold	1
π^c	Connection probability with distancing	0.5
h_1	Health demand parameter	0.55
h_2	Health demand parameter	0.1
π^m	Baseline death probability	0.02
h_3	Death probability parameter	0.06
π^r	Baseline detection probability	0.01
ψ	Adjustment of detection probability	0.0003
π_{max}^r	Upper bound of detection probability	0.1
$\mathcal{I}_0$	Number of initially infected	5

Table 8 contains the parameters used in the policy experiments.

Table 8: Policy parameters

Description	Value
Duration CIGS (months)	12
One-off liquidity injection (per firm)	10
Duration continuous liquidity support (months)	12
Duration credit guarantees (months)	12
Duration equity injections (months)	12

Table 8 – continued from previous page

Description	Value
Duration income support (months)	6
Size of income support (share of unemployment benefit)	0.5