Journal of Economic Behavior and Organization 217 (2024) 312–333

Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

journal homepage: www.elsevier.com/locate/jebo

Research Paper

Belief-driven dynamics in a behavioral SEIRD macroeconomic model with sceptics ☆

Christian R. Proaño [a,b,*], Jiri Kukacka [c,d], Tomasz Makarewicz [e]

[a] Otto-Friedrich-Universität Bamberg, Germany
[b] Centre for Applied Macroeconomic Analysis, Australian National University, Australia
[c] Charles University, Faculty of Social Sciences, Institute of Economic Studies, Czechia
[d] Czech Academy of Sciences, Institute of Information Theory and Automation, Czechia
[e] Bielefeld University, Germany

ARTICLE INFO

JEL classification:
E3
E7
I1

Keywords:
Endogenous cycles
COVID-19 pandemic
Bounded rationality
Heterogenous expectations

ABSTRACT

The reluctance of a non-trivial fraction of the population to adhere to social distancing measures – and even to get vaccinated – during the COVID-19 pandemic represented a challenge for imposed public health policies in many countries around the world. Against this background we study the impact of boundedly rational perceptions for the dynamics of epidemics such as the COVID-19 pandemic in a standard epidemic model extended by a stylized macroeconomic dimension similar to Atkeson et al. (2021). We illustrate through which channels misperceptions or even "scepticism" concerning the infectiousness of the disease or its mortality rate may undermine the effectiveness of lockdowns and other public health policies in the long-run.

1. Introduction

While the epidemiological dynamics of diseases such as COVID-19 have been modeled extensively, e.g., in the literature on compartmental epidemiological models based on the seminal work by Kermack and McKendrick (1927), the role of economic conditions in the evolution of such epidemics was much less investigated until recently, when works by Atkeson et al. (2021), Eichenbaum et al. (2021), and others led to an outright explosion of studies in the new field of "Pandenomics" (Cliffe, 2020).

While in frameworks such as Eichenbaum et al. (2021) the agents fully understand the epidemic and macroeconomic dynamics and consider the related health and economic risks in a rational manner, the reactions of a non-trivial number of individuals in many countries seem to suggest that other apparently "non-rational" factors may have played an important role in their behavior during the COVID-19 pandemic. In this context, one important and relatively less understood phenomenon of the COVID-19 pandemic was the existence of a non-trivial fraction of "sceptics" among the population, who can roughly be divided into two groups. The members of the first group either doubted the sheer existence of the COVID-19 virus (or explained its existence through conspiracy theories) or seemed to be generally reluctant to adhere to the social distance measures ordained by governments around the world. The second

☆ We would like to thank Augusto de la Torre, Blake LeBaron, Tiziana Assenza, Herbert Dawid, Giorgos Galanis, Alan Kirman, Leonardo Bargigli, Giorgio Ricchiuti, Benjamin Jeffries and other participants at the CEF 2021 conference and at the 1st DISEI Workshop on Heterogeneity, Evolution and Networks in Economics in Florence and at the Graz Workshop "Modeling the COVID-19 Pandemic" for helpful comments and suggestions on previous drafts of this paper, as well as the two referees of this journal. The usual disclaimers apply.

* Corresponding author at: Otto-Friedrich-Universität Bamberg, Germany.
E-mail address: christian.proano@uni-bamberg.de (C.R. Proaño).

https://doi.org/10.1016/j.jebo.2023.11.011
Received 29 August 2022; Received in revised form 12 October 2023; Accepted 8 November 2023

Available online 22 November 2023

group did acknowledge the danger of the COVID-19 pandemic, but believes them to be outweighed by the social and economic costs of lockdowns. These sceptics, known e.g. in Germany as "Querdenker" (lit. "lateral thinkers"), not only obtained a certain amount of prominence in the political discourse, but also compromised the long-run effectiveness of the vaccination campaigns by hindering the accomplishment of "herd immunity" through vaccination. Against this background, we address the following question: *what are the interactions and feedback loops between scepticism and efficiency of the governmental social distancing measures?*

Given the numerous and continuously increasing amount of studies that have emerged since the COVID outbreak by the end of 2019, a thorough survey is an impossible task to undertake. Nonetheless, there are a few studies worth mentioning. In the behavioral sphere, studies such as Eksin et al. (2019) and Di Guilmi et al. (2022) endogenize the reaction of susceptive individuals on physiological measures. Dasaratha (2020), for instance, extends a standard SIR model with endogenous meeting rates based on game-theoretic considerations, and Lux (2021) introduces social dynamics into the epidemiological process by assuming that the compliance of the population evolves adaptively according to the pandemic's perceived risk (though without taking into account the economic dimension as we do).[1] At the intersection between epidemiology and economics, Eichenbaum et al. (2021) were among the first to investigate the main economic transmission mechanisms of an epidemic such as COVID-19 in a model of forward-looking utility maximizing representative agents with rational expectations. Using the more elaborated heterogenous agents New Keynesian (HANK) approach, Kaplan et al. (2020) investigate the distributional consequences of social distancing measures in terms of income and wealth in a model calibrated for the United States. Using medium-scale agent-based models, Delli Gatti and Reissl (2020) and Basurto et al. (2023) analyze the joint dynamics of the epidemic and the economy in a model with various types of agents' heterogeneity. Finally, Basurto et al. (2020) study the efficiency of lockdown policies in a macroeconomic model with explicitly decentralized agent interaction.

Regarding the role of scepticism in the context of the COVID-19 pandemic, Allcott et al. (2020) documents the correlation between the political orientation of COVID scepticism and compliance to containment measures with the political orientation. Bursztyn et al. (2020) study the role of misinformation in mass media broadcasts and the adoption of preventative measures by the population, and Charron et al. (2023) study the impact of political polarization, see also Milosh et al. (2021). Mellacher (2023) also investigates the role of sceptics in the transmission of the disease in a theoretical model similar to ours, but in a purely epidemiological model with no interaction between the economic and the epidemiological spheres. To the best of our knowledge, Faia et al. (2022) is the only study besides ours that focuses on agents' perceptions of public health and economic or financial conditions, though using a different methodological approach. Using a survey experiment, the authors investigate how the ability of individuals to select public health- and economics-related information is related to their prior beliefs. However, they do not focus on the perceived trade-off between these two dimensions.

This recent but important strand of the literature considers the sceptic population as an exogenous factor. However, one can easily imagine that there exists a feedback loop between the sceptic fraction of the population and government policy. For instance, if the pandemic has already caused a significant number of infections, the public may be more willing to accept harsher lockdown restrictions on the economy, which in turn limits the reproduction rate of the virus and thus the future infectivity, as discussed e.g. by Lux (2021). In other words, the success of lockdown policies could, to some extent, be an *endogenous* process.

From a more theoretical perspective, the issue of scepticism is intrinsically related to the credibility of government policies. This issue has been thoroughly studied over the last three decades in a very different context, namely the credibility of monetary policy (see Lohmann, 1992 for an introduction and Aguiar et al., 2013 and Svensson, 2020 for two recent contributions). Among many interesting issues from this strand of literature, one that stands out for our research goal is the question of whether the public can trust the central bank to reach its policy target e.g. the inflation target. Usually, and following the seminal contribution by Barro and Gordon (1983), this question was addressed using models where all agents had rational expectations (RE) and thus fully understood the operating principle of the model. In recent times, and particularly after the 2007/08 Global Financial Crisis, the hegemony of the Rational Expectations Hypothesis has come under closer scrutiny. It has been increasingly acknowledged by the profession that it is important to focus on the effects of bounded rationality and learning, in particular to study models where the economic agents cannot simply move immediately to the RE solution, and instead have to learn how to form expectations based on past data (Assenza et al., 2021; Hommes and Lustenhouwer, 2019b). In such an environment of behavioral expectations, the credibility of the central bank can emerge endogenously under conditions that depend on the specific learning model. Conversely, if these conditions are not met, the public may learn to disregard the target of the central bank and to rely instead on off-equilibrium behavior, such as following inflation trends (Massaro, 2013; Mokhtarzadeh and Petersen, 2021).

The major theoretical challenge to the behavioral expectations literature is the so-called "wilderness of bounded rationality": there exist a large number of learning models that can have highly heterogeneous dynamic properties. One of the achievements of this literature is to limit this wilderness to a small number of candidates as a viable replacement of the RE, among which the two most prominent are Adaptive Learning (see Evans and Honkapohja, 2012 for a comprehensive introduction) and the Heuristic Switching Model (henceforth HSM; see Brock and Hommes, 1998 for an introduction and Hommes and Lustenhouwer, 2019a for a DSGE application).[2] In our paper, we will focus on the latter for two reasons. Firstly, this model is well supported by empirical and experimental studies in various financial and macroeconomic setups (see Kukacka and Sacht, 2023 for a recent example and

[1] See Funk et al. (2015) for a discussion of nine challenges in incorporating the dynamics of behavior in epidemiological models.

[2] In the behavioral literature, this modeling framework is also known as Adaptively Rational Expectations Dynamics (ARED) approach or discrete choice approach and encompasses models in which some fraction of the population can have fundamental or even rational expectations. See Matějka and McKay (2015) for the theoretical link between HSM and Rational Inattention models.

discussion). Secondly, as we will show in this paper, it is relatively straightforward to apply the HSM to the question of endogenous lockdown scepticism and much of our setup will directly reflect the previous studies on the credibility of monetary policy.

The HSM is based on the notion that agents, who need to forecast some economic variable and cannot directly calculate the RE solution, resort to a small number of simple rules of thumb, such as naive or trend-following expectations. The agents, however, are smart in choosing among them: they focus on the rules that were more successful in the past. In the case of monetary policy, if the central bank has so far kept the inflation close to (far away from) its target, agents will be more (less) likely to expect this fundamental inflation in the future as well, which in turn can bring the inflation closer to (farther away from) its target. This results in endogenous dynamics, in which the credibility of the central bank can emerge as a result of learning if the monetary policy is fine-tuned to the structure of the economy. As previously mentioned, the HSM can be easily adapted to the question of the COVID-19 scepticism. In our model, we will assume that the government imposes lockdown measures depending on the incidence rate, i.e. number of new daily infections per 100.000 inhabitants. However, the efficiency of the lockdown is scaled down by the share of sceptics among the general public, which itself is modeled with HSM. We will consider four versions of the model. In one dimension, the number of sceptics decreases with a higher number of either new daily infections or deaths.[3] In the second dimension, we will study sceptics who either do or do not react to the output gap.

The results of our analysis can be summarized as follows: First, the existence of an endogenous fraction of sceptics in the population, who do not follow government-ordered social distancing measures, can undermine the efficacy of such measures significantly. Second, the choice of the reference variables that determine the share of sceptics has important implications for the dynamics of the epidemic, and its economic consequences. In particular, the bigger the disconnect between the governmental and the public policy targets, the more unstable dynamics can emerge. Third, if the recovered population loses virus immunity over time, the model predicts that a weaker government response to the virus can lead to both a deeper recession and a higher death toll of the pandemic. Finally, a first empirical estimation attempt using a dataset combining the US economic and epidemic data provides statistically significant estimates of the two core behavioral model parameters. It thus suggests statistical evidence in favor of evolutionary switching dynamics towards scepticism and back towards trusting the government during the US pandemic of 2020 and 2021. To the best of our knowledge, our paper is the first one to highlight the trade-off between the economic and epidemiological impacts of COVID-19 and its role on the endogenous emergence of "sceptics" in a SEIRD-macroeconomic model. Thus, our study links the COVID-scepticism with the epi-macro literature discussed above.

The remainder of this paper is organized as follows: In section 2, we describe the basic epidemiologic SEIRD model and our baseline behavioral extensions. We then modify our model by incorporating sceptics into the model in section 3. Further, we investigate how the dynamics of the model change when the government and the public use alternative reference measures in section 4, and section 5 suggests potential policy extension of the model focused on capacity constraints at ICUs and two different immunity scenarios. Section 6 then completes the analysis with a Monte Carlo simulation study of the estimation potential of the model and subsequently applies the model empirically to US economic and epidemic data. Finally, we draw some conclusions from this paper in section 7.

2. A baseline reduced-form behavioral SEIRD-macroeconomic model

For the description of the epidemiological dynamics, we use a compartmental model which describes the evolution of a disease such as COVID-19 by dividing the total population into five categories: susceptible (S), exposed (E), infected (I) recovered (R) and deceased (D) similarly as in Kaplan et al. (2020).

In discrete time, the baseline SEIRD model reads

$$\Delta S_{t+1} = \gamma_S R_t - \beta_t S_t I_t / N_t \tag{1}$$

$$\Delta E_{t+1} = \beta_t S_t I_t / N_t - \sigma E_t \tag{2}$$

$$\Delta I_{t+1} = \sigma E_t - (\gamma_R + \gamma_D) I_t \tag{3}$$

$$\Delta R_{t+1} = \gamma_R I_t - \gamma_S R_t \tag{4}$$

$$\Delta D_{t+1} = \gamma_D I_t \tag{5}$$

$$N_{t+1} = N_t - D_t, \tag{6}$$

where N_t denotes the current population size, S_t the number of susceptible people at period t, E_t the number of exposed people, I_t is the number of infected people at t, R_t is the number of recovered people and D_t the number of deceased people. Please note that all the level variables are also subject to non-zero constraints. β_t is the so-called transmission rate, i.e. the number of extended contacts per day. The rate of mortality of the virus is denoted by γ_D, and the rate of recovery by γ_R, and $\gamma = \gamma_D + \gamma_R$ is the average

[3] The great uncertainty regarding the true mortality rate of the COVID-19 virus at the onset of the pandemic seems to have led many people to over- or underestimate its true threat. As discussed by a recent study by the World Health Organization (Msemburi et al., 2022), the estimated excess mortality due to the COVID-19 pandemic in the years 2020 and 2021 was about 14.83 million globally, 2.74 times more deaths than the 5.42 million previously reported.

Table 1
Baseline Calibration Parameters.

σ	incubation period	1/5.2	Atkeson (2020), Wang et al. (2020)
γ	Illness duration rate	1/18	Atkeson (2020)
γ_D	Death rate	0.004	Own parametrization
γ_R	Recovery rate	$\gamma - \gamma_D$	Own parametrization
γ_S	Reinfection rate	0.00	Own parametrization
α_y	Autoregressive output gap coefficient	0.95	Own parametrization
α_i	Infection rate impact on output gap	0.005	Own parametrization
α_b	Public health impact on output gap	0.04	Own parametrization
β_0	Baseline meeting rate	0.125	Wang et al. (2020)
$R_0 = \beta_0/\gamma$	Baseline transmission rate	2.25	Fauci et al. (2020)
μ_i	Intensity of choice parameter	2.5	Own parametrization
μ_y	Intensity of choice parameter	100	Own parametrization

duration of the disease. Finally, γ_S represents the waning of the viral immunity that recovered individuals initially acquire. In the special case of $\gamma_S = 0$, the immunity is permanent, while $\gamma_S > 0$ represents gradual immunity loss.[4]

While a purely epidemiological approach would consider the transmission rate to be as constant, a more realistic approach would take into account behavioral as well as policy-induced effects on this variable. In a similar vein as Atkeson et al. (2021) and Flaschel et al. (2022) and Lux (2021) we thus endogenize β_t as follows

$$\beta_t = \max\{0,\ \beta_0 - \phi_g(1-\omega_t)G_t\}, \tag{7}$$

where β_0 is the baseline transmission rate, $\phi_g G_t$ represents social distancing measures enforced by the government and ω_t is the share of sceptics, who undermine these measures. Accordingly, the higher this share is, the less efficient lockdown measures become.[5] The case with $\omega_t = 0$ represents a scenario with full governmental credibility (akin to the RE solution in monetary policy problems; we will discuss how $\omega_t \geqslant 0$ evolves endogenously in the next section).

As discussed by Eichenbaum et al. (2021), the two main channels through which economic activity is related and, in fact, influence the number of transmissions are (a) the amount of time spent by people at their working places with other workers and (b) the amount of their time spent in consuming (buying) goods and services.[6] As an epidemic affects an increasing fraction of the population, both types of activities may become less feasible, either because people are too sick to leave their homes, or because people decide to stay at home and avoid being infected voluntarily. Further, social distancing policies directly affecting the transmission rate β_t through the imposition of social distancing measures or eventual lockdowns also have a negative impact on economic activity. Accordingly, we specify aggregate economic activity (represented here by the output gap y_t) as follows:

$$y_t = \alpha_y y_{t-1} - \alpha_i I_t/N_t + \alpha_b(\beta_{t-1} - \beta_0) \tag{8}$$

where α_i represents the impact of I_t/N_t on y_t due to purely *physiological* factors, α_b the impact of the transmission rate on economic activity, and α_y the intrinsic persistence of the output gap process. It should be noted that when $I_t = 0$ and $\beta_t = \beta_0$ (no pandemic scenario), the output gap quickly converges to its natural steady state of $y_t = 0$.[7]

Unless otherwise stated, the following simulations are based on the parameter values reported in Table 1. For the epidemiological part of the model, we use parameter values used in standard studies such as Atkeson (2020), Wang et al. (2020), and Fauci et al. (2020). In particular, following Atkeson (2020) and Wang et al. (2020), we assume an incubation period of 5.2 days and an average duration of the disease of 18 days. Further, we assume for the basic reproduction ratio, defined as

$$\mathcal{R}_0 = \beta_0/\gamma,$$

the value of 2.25 proposed by Fauci et al. (2020).[8] Finally, the death rate of the virus seems to fluctuate around the 2% threshold, thus we assume $\gamma_D = 0.02/18 \approx 0.004$ (overall death rate divided by the illness duration).

In numerous countries a key measure for the evolution of the COVID-19 epidemic was the so-called incidence rate $\mathcal{I}_t$, namely the number of new cases σE_t per 100.000 persons.[9] In the following analysis, we will thus assume that this variable becomes the basis of the lockdown policy with $G_t = \mathcal{I}_{t-1}$. We have experimented with alternative policy specifications (namely $G_t = \Delta D_{t-1}$ – that

[4] In the Figures of section 3, we assume that $\gamma_S = 0$, and relax this assumption in section 4. Further, one can also introduce new variants of the virus with $\gamma_{S,t} = 0$ for most t and $\gamma_{S,\tau} = 1$ for some specific τ. We will leave this extension of the model for future studies.

[5] Note that this specification captures two alternative scenarios: sceptics simply refusing to follow the official guidelines, or sceptics putting political pressure on reducing the severity of these restrictions compared to the government's default policy.

[6] This is, of course, an oversimplification which is qualified by the capability of working from home (available, of course, only to a fraction of the working population), as well as by the use of delivery and takeaway services for the purchase of goods and to some extent, services.

[7] Given that the model has a relatively fast, daily frequency, we leave out standard economic factors, such as Philips Curve considerations, or the impact of expectations on savings.

[8] While Wang et al. (2020) propose a value of 3.1 for the description of the outbreak in the Chinese city of Wuhan, estimates for Western countries range between 2 and 3 (European Center for Disease Control). With a value of 2.25 we are thus on the lower end of the values proposed so far.

[9] In Germany, for instance, the 7-day average incidence rate has been prominently used as a threshold value for the implementation of sharp social distancing public policy measures.

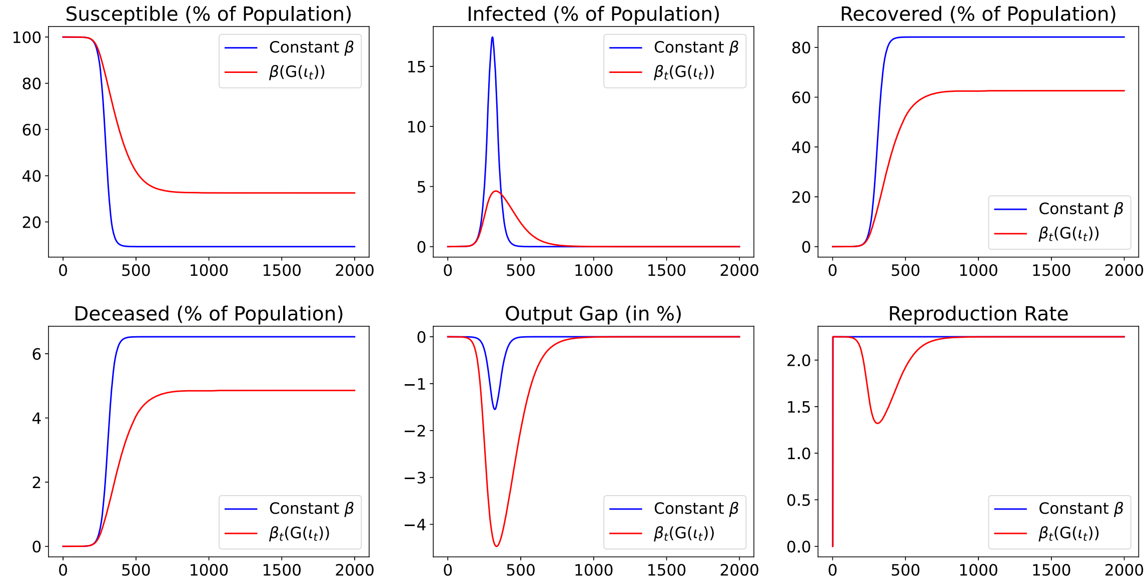

Fig. 1. The baseline behavioral SIR model with exogenous (baseline) and endogenous (scenario 1) number of extended contacts per day β_t following an initial infection of 100 persons of the population. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

means that the government reacts to the number of daily deaths – and a threshold policy in which higher levels of G_t are triggered by threshold levels of $\mathcal{I}_{t-1}$). Still, the model seemed to lead to similar qualitative policy implications. We leave this issue for future studies that could try to identify the optimal policy rule (see Basurto et al., 2023, for some discussion).

Fig. 1 illustrates the baseline dynamics of the model under $\beta_t = \beta_0$, i.e. without any social distancing measures imposed by the government, and thus a constant $\mathcal{R}_0$, and under the assumption of a varying β_t. The epidemiologic dynamics are well known by now not only to epidemiologists but to a broader public, including economists. An initial infection of a small fraction of the population leads ceteris paribus, and in particular under an unchanged reproduction rate $\mathcal{R}_0$ (resulting from an unchanged number of daily extended contacts $\beta_t = \beta_0$ and represented by the baseline blue lines in the individual graphs in Fig. 1) to a rapid infection of an increasing number of susceptible individuals (which we assume are the totality of the population). Without any public policy aimed at reducing the transmission rate β_t and thus of the reproduction rate $\mathcal{R}_t$, the epidemic enters an exponential growth phase which leads to a swift spread of the disease over the population in a short period of time. Given the (still constant) mortality rate assumed so far, the rapid increase in the number of infected also leads to a significant number of deceased persons. Since the rate of infection (number of infected people relative to the population) negatively affects the level of economic activity – represented in our model by the output gap variable, see equation (8) – the spread of the disease leads to a negative output gap. Given our parametrization, this impact is relatively small, leading, at the peak of the infection, to about a 1% decrease in economic activity due to purely physiological reasons, i.e. solely due to the sickness-related reduction of production in the economy.

When a social distancing policy is implemented, the current transmission and reproduction rates β_t and $\mathcal{R}_t$ respectively, are reduced. This has opposing effects on the epidemic and economic spheres: In the epidemic sphere, the containment policy leads to a significant reduction in the number of newly infected people per day and, given the constant mortality rate of the epidemic, also to a lower number of deceased people in the long run. In the economic sphere, by contrast, the slump in economic activity, which in the previous case was only due to the pure physiological effects of the epidemic, is magnified by the containment policies.

3. Incorporating sceptics

So far, we have assumed that the contact rate β_t and the reproduction rate R_t^0 vary only due to government-imposed social distancing measures. Further, we have also implicitly assumed that these measures are followed and accepted by the totality of the population, which leads to successful containment of the disease, as discussed in the previous section. However, as we have discussed before, a non-trivial phenomenon undermined the success of public health policies. It even threatened the achievement of herd immunity in the long-run (such as is the case in Germany with its rather low rates of vaccinated people despite the vast vaccine supply) was the emergence of COVID-19 "sceptics", who did not adhere to social-distancing measures.[10] In the following, we model the compliance to social-distancing measures as an endogenous choice based on the perceived health-related and economic costs using the Heuristic Switching framework.

[10] Fortunately, the COVID-19 virus mutated through 2022 to a variant with a much lower mortality rate that eased the constraint created by this "sceptic" behavior. We may not be so lucky next time.

3.1. A heuristic switching model of scepticism

As discussed e.g. by Hommes (2013), the Heuristic Switching Model (HSM) is based on the following idea. Suppose that in a dynamic system represented by a vector of state variables z_t, agents need to forecast certain variables $x_{t+1} \in z_{t+1}$. For some reason, the rational solution is unavailable to or cannot be computed by the agents. Instead, they rely on a set H of heuristics of a form $x^e_{h,t+1} = H_h(z^t)$, where the superscript t denotes the history of the system z_t (including past values of x_t). Heuristics $H_h(\cdot)$ tend to be simple, for instance, a simple naive rule is given by $x^e_{t+1} = x_t$ and a simple bias can be represented by $x^e_{t+1} = b$ for some bias value b.

Even though the rules are simple, the agents are smart in choosing them. They focus on some measure of success $U_h(t)$ for each heuristic, and in the context of forecasting financial or macro variables, a popular metric is squared forecasting error with $U_h(t) = -(x^e_{h,t} - x_t)^2$. Hence, the agents switch towards the heuristics that have been more successful. Define the logit transformation

$$\omega^h_t = \frac{\exp(\mu U_h(t))}{\sum_{k\in H} \exp(\mu U_k(t))} \tag{9}$$

as the share of the population that at period t will use heuristic h. Parameter μ represents the intensity of choice of the agents: when $\mu = 0$, agents choose heuristics at random, when $\mu \to \infty$, they immediately switch to the best heuristic, and for finite positive values of μ, switching is more gradual.

Our model of scepticism will have almost the same structure. Instead of focusing on forecasting heuristics, we will assume that there are two forces that agents are subject to:

Pull forces: Agents are *pulled* towards trusting the government – and thus to reduce their scepticism – if they observe that the pandemic "is real". In this context, we will consider two specifications[11]:
- $U^I_t = \Delta I_t$ – the number of newly infected persons;
- $U^D_t = \Delta D_t$ – the number of newly deceased persons.

Note that both specifications (with U^C_t for $C \in \{I, D\}$, C standing for credibility) refer to observable variables and thus imply a tangible experience of the pandemic and its current severity.

Push force: Agents are *pushed* towards scepticism for two reasons:
- $U^N_t = 0$ represents the case in which people are inclined to simply disbelieve the pandemic (or its severity);
- $U^Y_t = -y_t$ represents the case in which people worry about the economic consequences of the social distancing measures.

The two push forces $S \in \{N, Y\}$ (S standing for scepticism) represent the two types of resistance towards lockdown measures discussed in the Introduction.

Accordingly, the share of sceptics is endogenized as

$$\omega_t = \frac{\exp\left(\mu_y U^S_t\right)}{3 + \exp\left(\mu_y U^S_t\right) + \exp\left(\mu_i U^C_t\right)}, \tag{10}$$

for $S \in \{N, Y\}$ and $C \in \{I, D\}$. Note that this leads to four model specifications, which we will refer to as I/N, D/N, I/Y and D/Y models. We will consider the I/Y specification as the benchmark model since it is likely the most realistic one.

In the standard HSM, the heuristics forecast the same variable, such as inflation. In contrast, our agents weigh two different factors, namely the current occurrence of infection (expressed in the count of either new infections or deaths per 100.000 people) and the state of the economy. To bring these two variables to a common factor, we weigh them with two different intensity of choice parameters. Secondly, the number of infections and deaths can quickly start to grow exponentially. On the other hand, a forecasting heuristic can make both a positive and negative mistake, whereas our push and pull forces are always non-negative, with $\Delta \mathcal{I}_t, \Delta D_t, -y_t \geqslant 0$. Therefore, for the sake of numerical stability and easier interpretation, we settled on a model, which is analogous to an HSM with (1) naive forecasting rules and (2) absolute, instead of squared, forecasting errors.

Finally, note that for $y_{t-1} = 0$ and, depending on the model specification, $\Delta I_t = 0$ or $\Delta D_t = 0$, the share of sceptics becomes $\omega^o = 1/(3+1+1) = 0.2$, i.e. by default 20% of the population does not support the government's containment policies. Assuming such a high population of sceptics in the steady state may seem quite exaggerated. However, as Fig. 2 illustrates, the share of the population disagreeing that vaccines (in general, not only COVID-19 vaccines) are safe or effective goes up to 20-30% in some countries like France.

3.2. Micro-foundations of the heuristic switching model

In the standard versions of HSM, the popularity of forecasting heuristics (as defined by equation (9)) depends on their relative predictive performance. Our model of "pull-push" forces has a similar structure, where the forces represent the attitude of households towards social distancing measures: whether the households should *trust* or *distrust* the government and its policy. Below we provide a stylized choice model that can serve as micro-foundations for the HSM from the previous section.

[11] In both cases, we consider *absolute numbers* here as media often reports epidemic numbers in absolute and not relative terms, for example, relative to the overall population or relative to the number of people currently infected. Remark that unless the population dramatically decreases due to the high death toll of the pandemic, the overall and relative number of infections are proportional and thus have no bearing on the model up to the rescaling of the μ parameters.

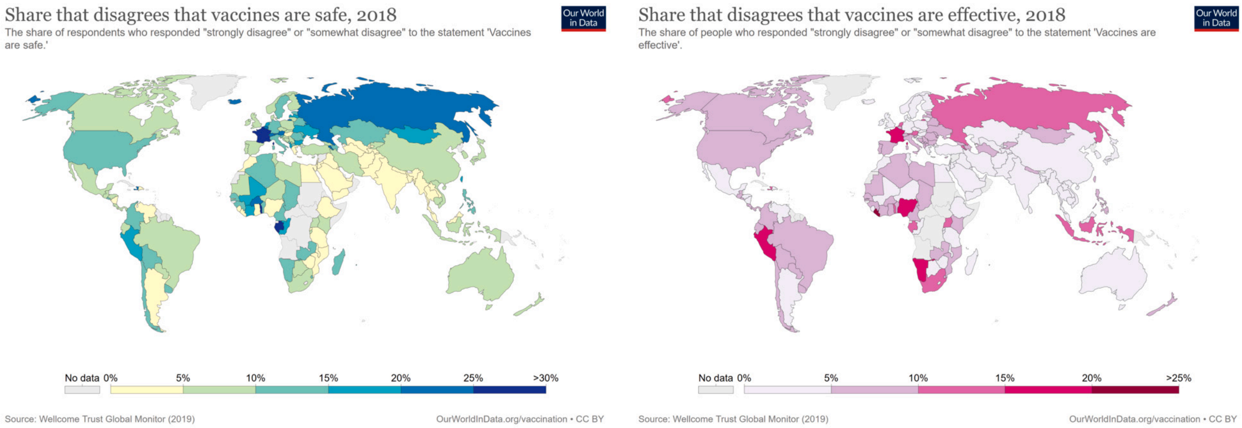

Fig. 2. Vaccine scepticism around the world. Source: Welcome Trust Global Monitor 2018 through Our World in Data.

Consider a mass of symmetric households, which during a pandemic have to decide on how to respond to the announced social distancing measures G_t. Given the sharp uncertainty about the dynamics of the pandemic, as well as the high-frequency nature of the issue (relative to the more standard quarterly macro setup of DSGE models), we assume that the households are subject to cognitive constraints on their rationality and thus address the problem in a myopic fashion.

Suppose that in period t, household i sets its compliance level to $K_{i,t} \geqslant 0$. The household receives utility according to the function $U_{i,t}(K_{i,t}) = \ln u_{i,t}(K_{i,t})$, where

$$u_{i,t}(K_{i,t}) = \kappa_1 K_{i,t} \exp(Z_t) + \kappa_2 K_{i,t} \exp\left(\kappa_3 y_t\right) - \kappa_4 \left(K_{i,t} - G_t\right), \tag{11}$$

for some $\kappa_1, \kappa_2, \kappa_3, \kappa_4 \geqslant 0$.[12]

This utility function represents three separate effects of the lockdown compliance. Firstly, while the overall well-being of an agent decreases with hazard Z_t, individual compliance $K_{i,t}$ mitigates the epidemiological hazard by κ_1 for a given hazard Z_t, increasing their well-being. This becomes clear when one considers the partial derivative of the agent's utility of $K_{i,t}$ related to Z_t, $\kappa_1 \exp(Z_t)$, which is increasing in Z_t. Secondly, extensive lockdowns require shutting down some sections of the economy, which can affect the personal income of many households, particularly those that cannot work from home or are at risk of an extended spell of unemployment. This vulnerability on the aggregate level is captured by the parameters κ_2 and κ_3. Finally, some may find the social distancing to be personally challenging, due to psychological issues (like mental health risks from extended isolation) or political/ideological concerns (like a mistrust towards governmental intervention). This is measured by the parameter κ_4 and we assume that the government's target measure G_t plays an additional role as a reference point to this effect (for instance by normalizing what is the "new normal").

Under Rational Expectations, the agents would optimize (11), taking into account the fact that Z_t is in fact a function of the present and past levels of K_t. Instead, suppose that due to cognitive or informational limitations agents consider two basic heuristics:

- trust in government and comply with $K_{i,t} = G_t$, which yields utility equal to

$$U_t^C = \ln\{\kappa_1 G_t \exp(Z_t) + \kappa_2 G_t \exp\left(\kappa_3 y_t\right)\};$$

- remain sceptic with $K_{i,t} = 0$, which yields utility equal to $U_t^S = \ln\{\kappa_4 G_t\}$.

Assuming the simplest switching mechanism (without memory and with the intensity of choice set to 1), the share of sceptics at period t becomes

$$\begin{aligned}\omega_t &= \frac{\exp\{U_t^S\}}{\exp\{U_t^S\} + \exp\{U_t^C\}} = \frac{\kappa_4 G_t}{\kappa_4 G_t + \kappa_1 G_t \exp(Z_t) + \kappa_2 G_t \exp\left(\kappa_3 y_t\right)} \\ &= \frac{\exp(-\kappa_3 y_t)}{\kappa_2/\kappa_4 + (\kappa_1/\kappa_4)\exp\{\kappa_3((Z_t/\kappa_3) - y_t)\} + \exp(-\kappa_3 y_t)}.\end{aligned} \tag{12}$$

What is the hazard measure Z_t? Remark that the risk of catching the disease increases 1) proportionally with the overall economic activity and 2) exponentially with the size of the infected population (and is thus correlated with the lagging indicator of deaths). We assume that our boundedly rational agents therefore use either $Z_t = \Delta I_t + \kappa_3 y_t$ or $Z_t = \Delta D_t + \kappa_3 y_t$. Without any loss of generality we

[12] In the HSM literature, it is typically assumed that the agent heterogeneity stems from the relative popularity of forecasting heuristics, whereas the agents are otherwise symmetric. For the sake of simplicity we will follow this approach in our model.

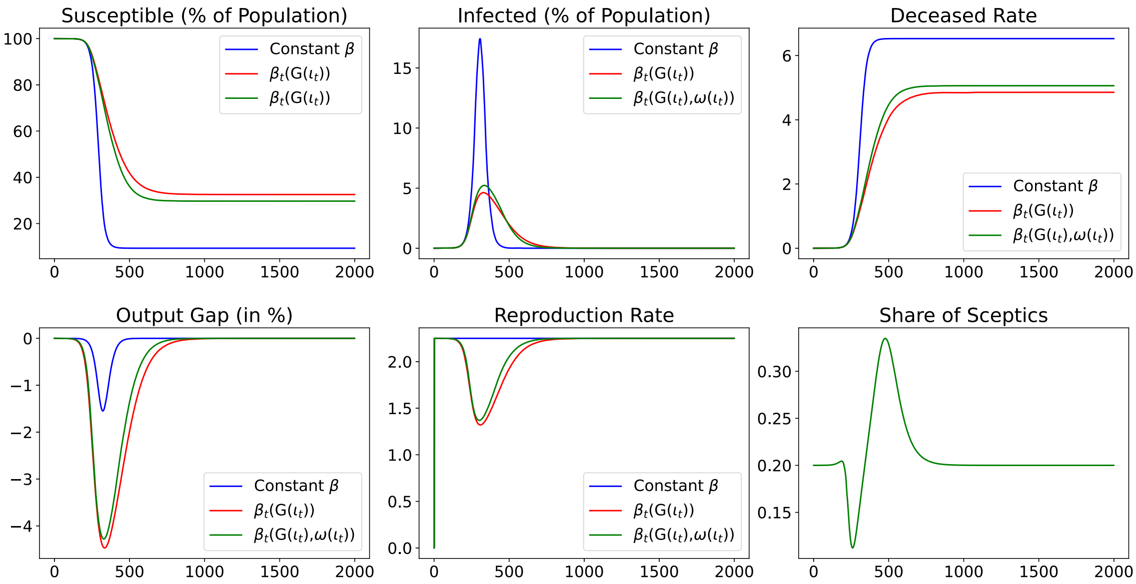

Fig. 3. The SEIR model with endogenous transmission rate β_t and a varying fraction of sceptics in the population with $\phi_g = 0.025$ and $\mu_y = 100$, $\mu_i = 2.5$.

can further assume that $\kappa_3 = \mu_y = \mu_i$, $\kappa_2/\kappa_4 = 3$ and $\kappa_1/\kappa_4 = 1$. With these assumptions put together, the sceptic share (12) coincides with the one from the main model if the agents consider y_t in their decision problem (model variants I/Y and D/Y).

The share (12) has one interesting special case for $\kappa_3 = 0$, when the agents are indifferent to losses in the output gap. Here ω_t can be rewritten as

$$\omega_t = \frac{\exp(0)}{\kappa_2/\kappa_4 + (\kappa_1/\kappa_4)\exp(Z_t) + \exp(0)}. \tag{13}$$

We take $\kappa_2/\kappa_4 = 3$ and $\kappa_1/\kappa_4 = 1$ as before. Furthermore, we simplify the hazard measure to exclude the output gap as well, with $Z_t = \Delta I_t$ or $Z_t = \Delta D_t$, which yields the I/N and D/N variants of our model. The pandemic demonstrated that this case is highly unrealistic. On the other hand, some countries tried to incorporate extensive measures to "freeze" the labor market (such as the aforementioned German Kurzarbeit), in an attempt of alleviating the negative economic consequences of lockdowns on the households. It is, therefore, important to study this case as an ideal policy target.

3.3. Example dynamics

Fig. 3 illustrates the dynamics of the I/Y model for $\phi_g = 0.025$ and $\mu_y = 100$, $\mu_i = 2.5$. As it can be observed, for the chosen parameter constellation, the share of sceptics in the population, after a very slight initial increase, decreases from 20% down to nearly 10% in the initial phase of the epidemic, where the number of new infections increases significantly, and the output gap decreases. However, as the latter increases even further and the percentage share of infected persons in the population decreases, the popular sentiment changes, and the share of sceptics increases, leading to a slight increase in the reproduction rate relative to the baseline scenario where no sceptics were present, i.e. where the transmission rate (and by extension the reproduction rate) depended solely on the government's containment policies ($\beta_t(G_t)$).

When public health is relatively less valued by the population ($\mu_i/\mu_y = 0.015$), in comparison to the previous parametrization ($\mu_i/\mu_y = 0.025$), the share of sceptics increases in relation to the previous case (which we depict here again for better comparison), as is illustrated in Fig. 4. The smaller reduction in the reproduction rate resulting from the compromised effectiveness of the government's containment policies due to the sceptics has non-trivial results for this second parametrization. Regarding the economic sphere, the sceptics' non-adherence to the social distancing measures reduces the fall in the output gap to some extent, as can be clearly observed. However, this comes at a great cost, as the disease's mortality rate (deceased persons as a percent of the total population) in the long-run increases by about 1.5 percentage points, from about 11% to about 12.5%.

4. Varieties of scepticism and policy trade-offs

In this section, we first investigate how stricter policy measures may interact with endogenous scepticism in the I/Y model, and then discuss the trade-offs that policy-makers are faced under the I/N, the D/N and the D/Y models besides the I/Y model.[13]

[13] In the main part of this paper we will focus on the results for the model with the intensity of choice parameters set to $\mu_i = 1$ and $\mu_y = 1$. A bifurcation analysis suggests that the model has largely robust dynamics to this specification. As intuition would suggest, the higher (lower) μ_i compared with μ_y, the weaker (stronger)

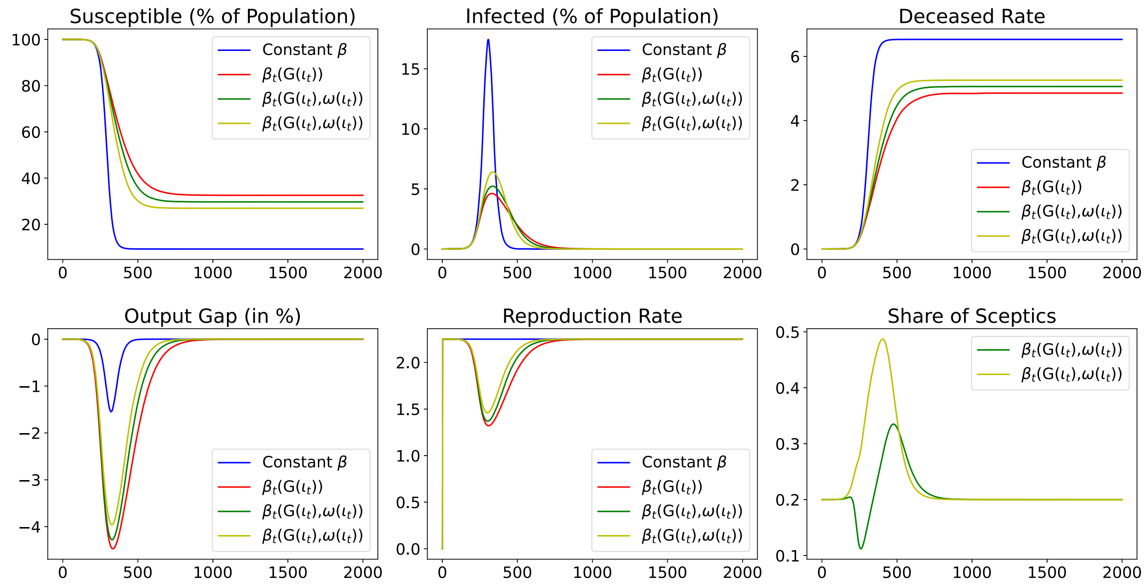

Fig. 4. The SEIR model with endogenous transmission rate β_t and a varying fraction of sceptics in the population with $\phi_g = 0.025$ and $\mu_y = 100$, $\mu_i = 2.5$ (green lines) and $\mu_i = 1.5$ (yellow lines).

4.1. *Stringency of policy and scepticism*

Can a more stringent social distancing policy curb scepticism? We consider three levels of governmental reaction with $\phi_g \in \{0.5, 2.5, 50\}$ in the I/Y model, which henceforth we will refer to as weak, medium, and strong reactions, respectively.

As illustrated in Fig. 5, a weak government reaction implies a relatively small drop of the transmission rate β_t in the "burn-through" phase of the pandemic, thus contributing to the spread of the pandemic. On the other hand, the number of new daily infections is so high, that even with a severe drop in the output, the share of sceptics drops to zero, hence the weak response does still have some impact on the transmission rate. This is not enough to stop the burn-through from occurring, however, and once everyone gets infected and dies or recovers, the output gap goes back to zero, and the share of sceptics converges to its default value of 20%.

Very different dynamics occur for medium and strong government reactions. In both cases, the government reacts much more stringently and quickly to the rise in new infections, which causes a visibly stronger drop of the transmission rate, to around a third of its pre-pandemic level – and thus a prolonged recession of around 5%. However, despite similar economic outcomes, the ratio of sceptics is very different in these two scenarios: While it settles close to 20% under the medium policy, under the strong policy, the great majority of the population (90%) become sceptical. We observe the following feedback mechanism. A more stringent or stronger response means that the government is able to decrease the number of new infections much more quickly. This, in turn, implies that the households, in comparison to a medium reaction scenario, observe a similar drop in the output gap, but with a substantially smaller number of new infections, and the latter factor contributes to a higher level of scepticism. On the other hand, the higher levels of scepticism are counter-balanced by the stronger reaction of the government, and it seems that the net effect on the output gap settles at an equilibrium level, which does not depend on the particular strength of the government's reaction. The only difference is a slight change to the number of infections and thus sceptics. Finally, notice that the strong policy has one important side-effect: the government reacts with force to a higher incidence rate and thus the number of newly infected quickly goes down, which allows for a lower lockdown intensity (as measured through the meeting rate β_t), but at a price of a renewed wave of infections in a few periods later. The reason behind this oscillatory pattern of "flash lockdowns" is the length of the incubation period, which allows the virus to spread without being first detected from the incidence rate.

This result points to another interesting outcome of the model. Once the government decides to actually react to the spread of the epidemic, the specific speed of the response (with a relatively weaker or stronger reaction parameter) does not matter in the sense that weaker responses lead to higher infection numbers, which then forces the government to act anyway, only later. In other words, the government is not facing a trade-off between less or more severe lockdown, but between imposing a strict lockdown sooner or later.

government policy ϕ_g is required to achieve the same outcome. However, no additional dynamics occur or disappear for reasonable values of these two parameters, which only rescale the efficacy of a given government strength. See Appendix A for some related simulation results.

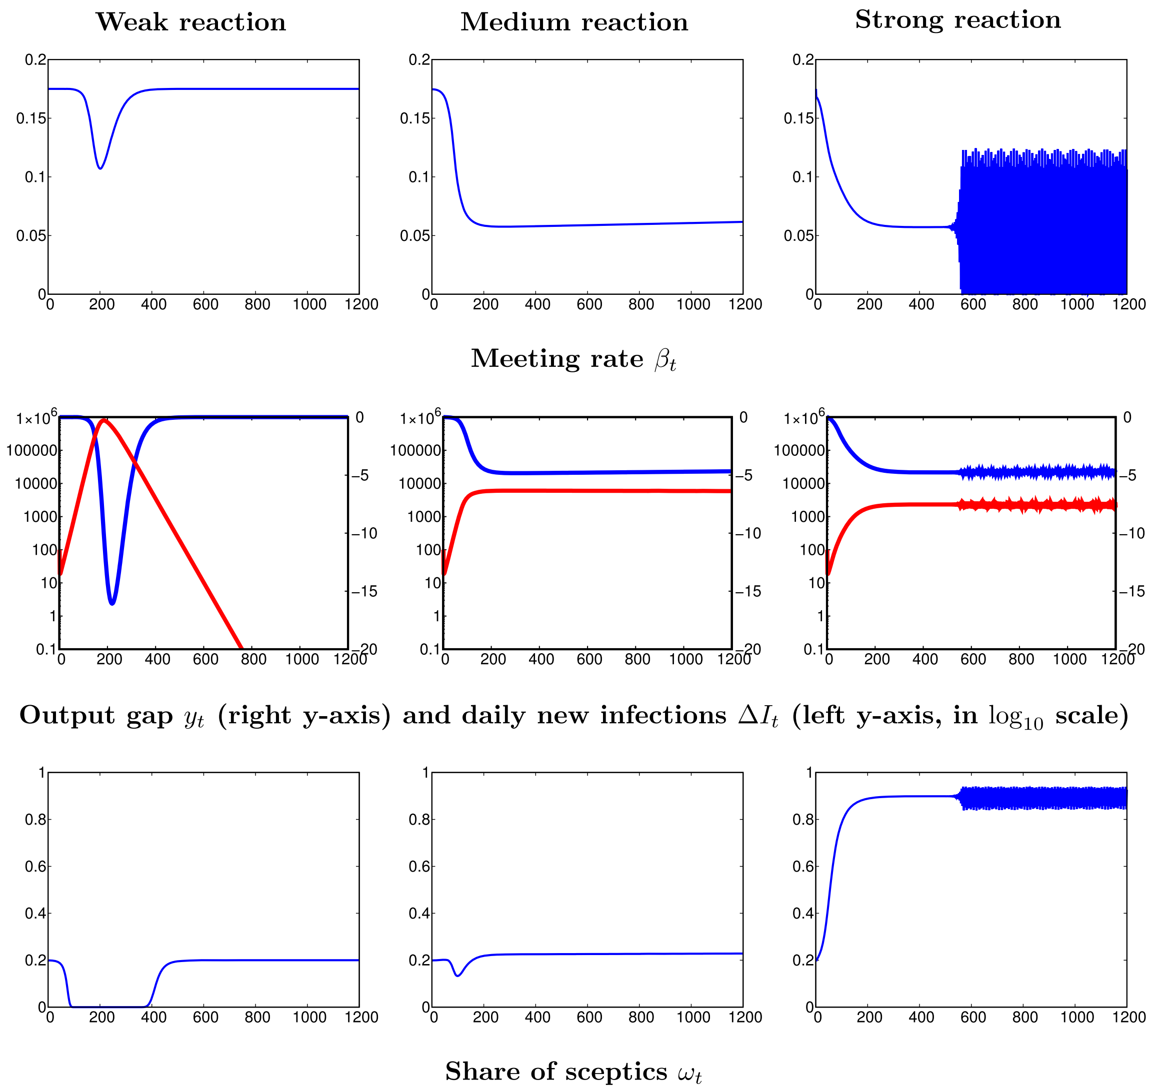

Fig. 5. ***I*/*Y* model:** Sample simulation for weak (left panels), medium (middle panels) and strong (right panels) reaction variants of the model, with ϕ_g equal to respectively 0.5, 2.5 and 50.

4.2. *Varieties of scepticism*

So far, we have considered the benchmark I/Y model, in which the share of scepticism reacts to the number of infections. Two questions can be raised: what if the share of sceptics decreases with the number of deaths, instead of new daily infections; and what if instead of the economic costs of lockdowns, the scepticism was fueled by a mere distrust towards the existence of the pandemic? To address this issue, we will consider the four versions of the model: I/N, I/Y, D/N and D/Y.

Sample simulations for a moderate government policy can be found on Fig. 6. The clear pattern is that three models, the I/N, I/Y and D/N have virtually indistinguishable dynamics. In other words, suppose that initially, people are sceptical of the existence of the virus, but gain more confidence in the government policy as the number of new daily infections increases – then it does not matter whether they suddenly look at the number of deaths instead of infections (the pull towards trusting the government) or their scepticism is now driven by the output gap considerations (the push towards scepticism). However, if people start caring about deaths (instead of new infections) and stop caring about the output gap *at the same time*, the feedback between these two forces leads to a breakdown of the efficiency of the public policy, and the government starts to struggle with imposing the lockdown rules.

Fig. 7 illustrates the mechanism behind these dynamics. For the two models I/N and D/N, the number of infections and deaths quickly reaches the point, where scepticism drops to its default level of 20%, and the government can run its lockdown policy quite efficiently. The same happens for the I/Y model, in which the output gap initially leads to high shares of scepticism, but the recession is quickly overshadowed by the toll of the pandemic. However, these dynamics are not present in the D/Y model.

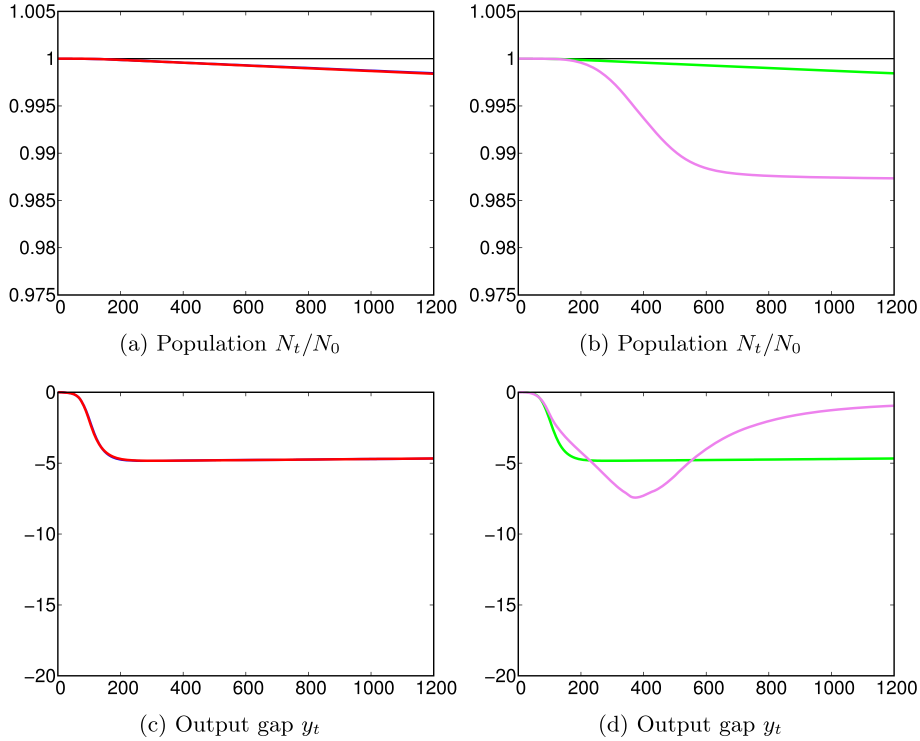

Fig. 6. Sample simulations for the four types of scepticism: I/N (blue lines), I/Y (red lines), D/N (green lines) and D/Y (violet lines) models, for medium government policy with $\phi_g = 2.5$. The top and bottom panels display the population and output gap, respectively.

In the D/N and D/Y models, as in the I/N and I/Y variants, once the number of infections starts to go up, the government tries to impose lockdown measures. However, unlike in the I/N and I/Y model, the households look at the number of deaths, which lag behind the incidence rate, hence they start as being highly sceptical (with $\omega_t \approx 1$). The difference between the D/N and D/Y models is that once the number of daily deaths hits a sufficiently high threshold, in the D/N variant, the households start to trust the government on the severity of the pandemic. On the other hand, in the D/Y model, this increase in deaths is offset by the output gap, keeping scepticism at a high level and thus rendering the government policy less efficient. This lets the virus go unhindered among the population, and by the periods $t = 400 - 600$, already half of the population has suffered an infection, fueling the recession. In other words, the effects of scepticism are paradoxical: because the population fears a recession, its opposition towards the lockdown measures renders those measures inefficient, fueling a burn-through of the virus through the population and a snap recession (as sick people stay at home for some time).

Once the virus has burned through half of the population, the number of daily infections starts to decline, which enables the government to loosen up the lockdown measures. These two factors contribute to economic recovery, a higher transmission rate, and lower rates of scepticism. Summing up, the dynamics in the D/Y model resemble the I/Y variant, but with weak policy parametrization $\phi_g \approx 0$, where a relatively quicker recession corresponds to the government's failure to contain the virus.

Fig. 8 displays the bifurcation diagrams for the population and the output gap, for the I/Y and D/Y models. It is clear that the difference between these two models, which we observed in the sample simulation for the medium strength government response $\phi_g = 2.5$, holds for a broad range of policy parametrizations. Remark that the model was calibrated to result in a 2% death toll of the infected population. This means that a simulation, by which end only 0.98 of the initial population remains alive, represents the epidemiologically worst outcome. On the other hand, the D/Y model simulations, in fact, yield results that are quite close to this scenario, even for a relatively strong governmental response, which allows the government to avoid this death toll in the I/Y model.

How to interpret these results? In our model, the government is always focused on the incidence rate as the proxy for the severity of the pandemic. On the other hand, scepticism exemplifies a lack of coordination between the government and the population in terms of policy goals. In the case of the I/Y and D/Y models, the bone of disagreement is that the public works on a "wider" preference structure, with economic consequences of the recession factored in next to the pandemic. In the D/N and D/Y models, the general public measures the severity of the pandemic with a "slower" indicator (number of deaths, which lag behind the infections). Our model suggests that as long as the dis-coordination between the government and the population is sufficiently small, the government can mitigate it through a stronger policy. However, once the *structural* reasons for the dis-coordination start to pile up, this becomes more and more difficult, and in D/Y model, becomes virtually impossible.

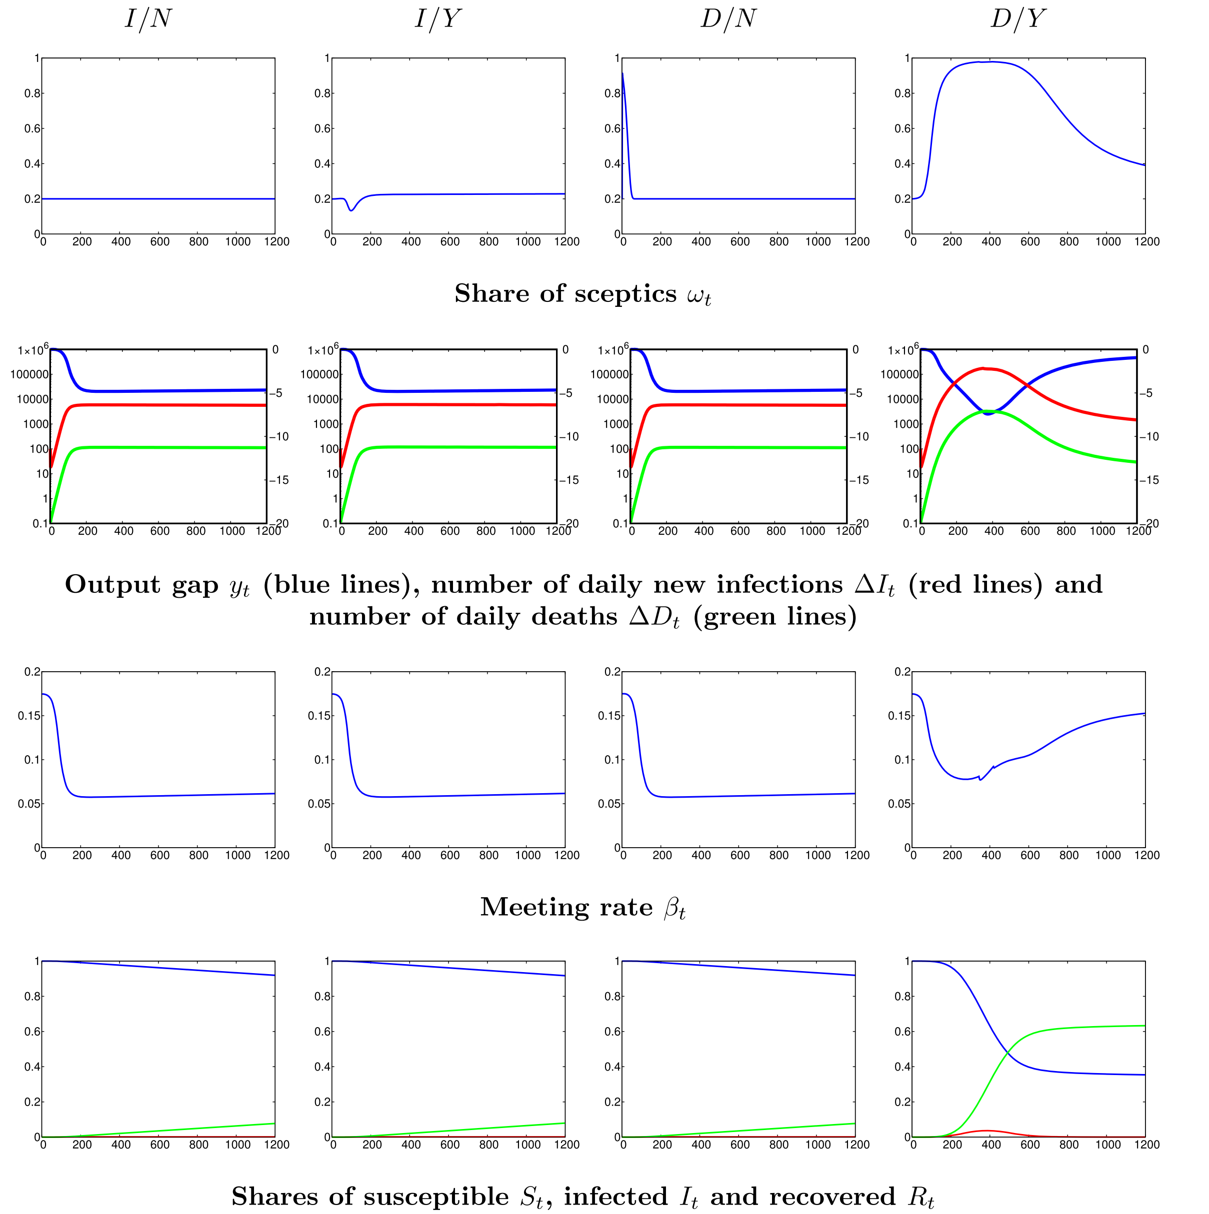

Fig. 7. Sample simulations for the four types of scepticism: I/N (left-most panels), I/Y (second-left panels), D/N (second-right panels) and N/Y (right-most panels) models, for medium government policy with $\phi_g = 2.5$. Remark that the new daily infections and deaths are presented on a $\log_{10}$ scale.

5. Policy extension: ICU capacity constraint and depleting virus immunity

In the previous sections, we studied the model under two critical assumptions: (1) recovered individuals obtain permanent immunity against the virus and (2) the death rate is constant across time. Both assumptions are, in fact, questionable. Firstly, recent studies suggest that the immunity against SARS-CoV-2 wanes over time (see for example Townsend et al., 2021). Secondly, the average mortality rate of this virus is, in fact, close to 2%, but it has been changing across time and between countries. This can at least to some extent be explained by the capacity of the healthcare system, which in most countries was not designed to handle a pandemic crisis at the scale caused by SARS-CoV-2.[14] From the outset of the pandemic, part of the motivation behind public lockdown measures was the fear that clogged hospitals and overburdened medical staff would be unable to attend all patients, resulting in excess deaths – and we now have many examples, that this fear was in fact justified, see Rossman et al. (2021) for an Israeli case study.

[14] Another factor that we leave out of the model is that the infectivity seems to be weather-dependent. This could be modeled with a time-varying $\sigma_t = \sigma_0 + \sigma_1 \sin(\sigma_2 t)$.

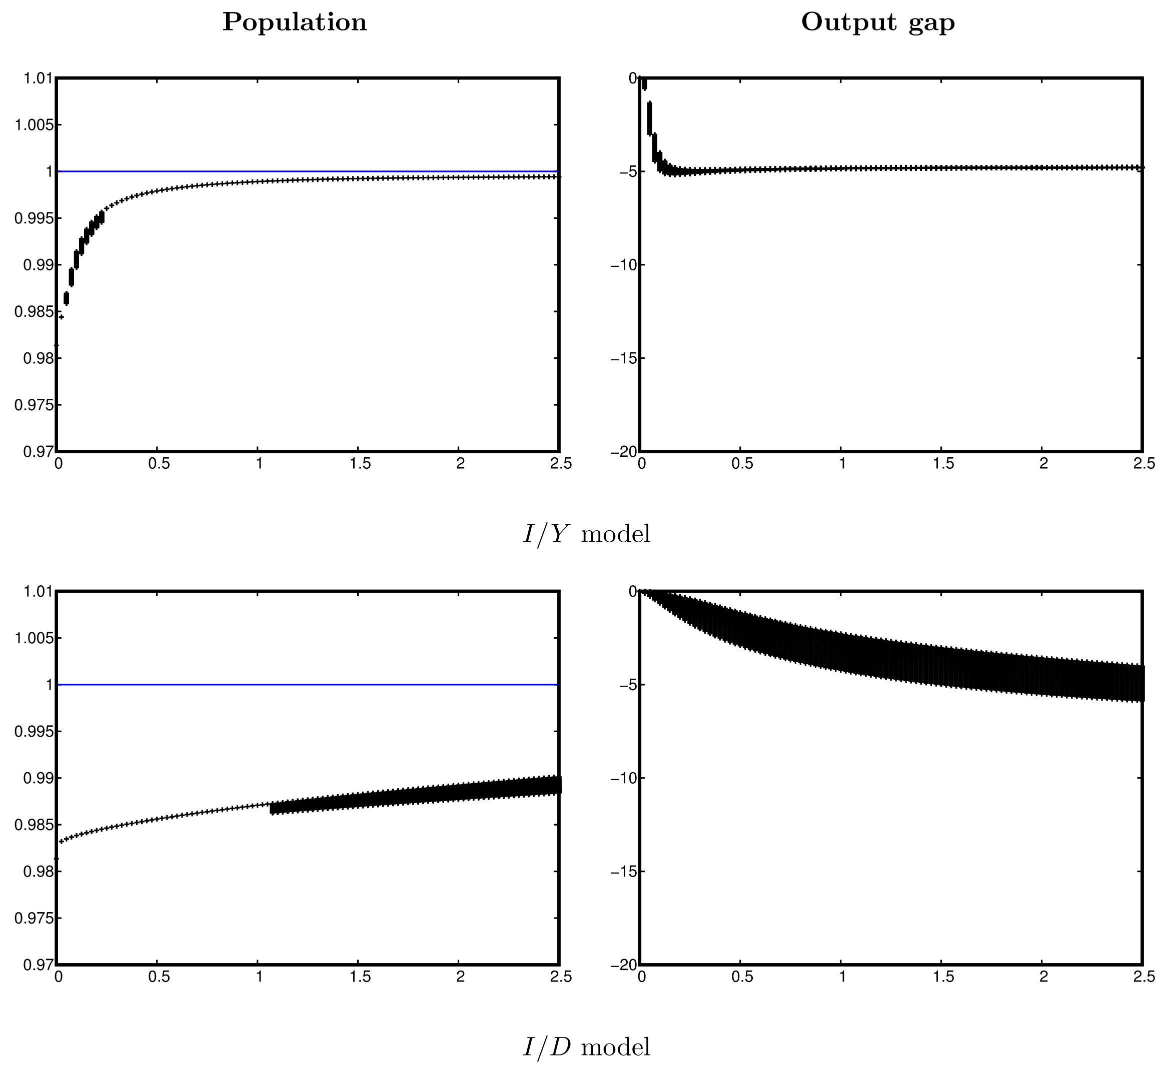

Fig. 8. I/Y **versus** D/Y **model:** Bifurcation diagrams of Population (left panels) and output gap (right panels) as functions of ϕ_g. Each vertical slice represents the range of outcomes in periods $t \in \{1101, \ldots, 1200\}$ for the given parameter ϕ_g.

In this section, we will study the effect of these two issues in our model. For the sake of tractability of the results, we will always consider the benchmark I/Y model, and we will focus on four variants resulting from the mix of the following two dimensions: Along the first dimension, we consider the cases where $\gamma_S = 0$ and $\gamma_S = 0.005$. These two cases represent a permanent immunity scenario and the scenario, in which the half-life of the virus immunity is around 140 days, respectively. In the second dimension, we consider a fixed mortality rate $\gamma_{D,t} = \gamma_D$ and what we call an "explosive mortality scenario", in which

$$\gamma_{D,t} = \begin{cases} \gamma_D & \text{if } I_{t-1} \leqslant \hat{I}, \\ \gamma_D + (\gamma_D^{max} - \gamma_D)\frac{I_t - \hat{I}}{N_t - \hat{I}} \in [\gamma_D, \gamma_D^{max}] & \text{else.} \end{cases} \tag{14}$$

Recall that $\gamma_D = 1/900$ is calibrated so that over the typical illness period of 18 days, 2% of infected will die. We set $\gamma_D^{max} = 1/180$, i.e. five times that mortality rate and $\hat{I} = 350000$ (tenfold of ICU beds in Germany). Equation (14) has the following interpretation. Suppose that 10% of the infected require hospitalization with an ICU bed.[15] As long as the number of currently infected people does not exceed $\hat{I}$, all COVID-19 patients receive the best possible medical care. However, once the number of infected passes this threshold, the hospitals start to get clogged, which results in excess deaths. This effect depends linearly on the number of "excess

[15] In practice, this number could be lower, however, we leave out second-order effects of peaks in COVID-19 cases. In particular, a large number of COVID-19 patients may result in patients with other ailments being "crowded out" from, or being afraid to use the hospitals.

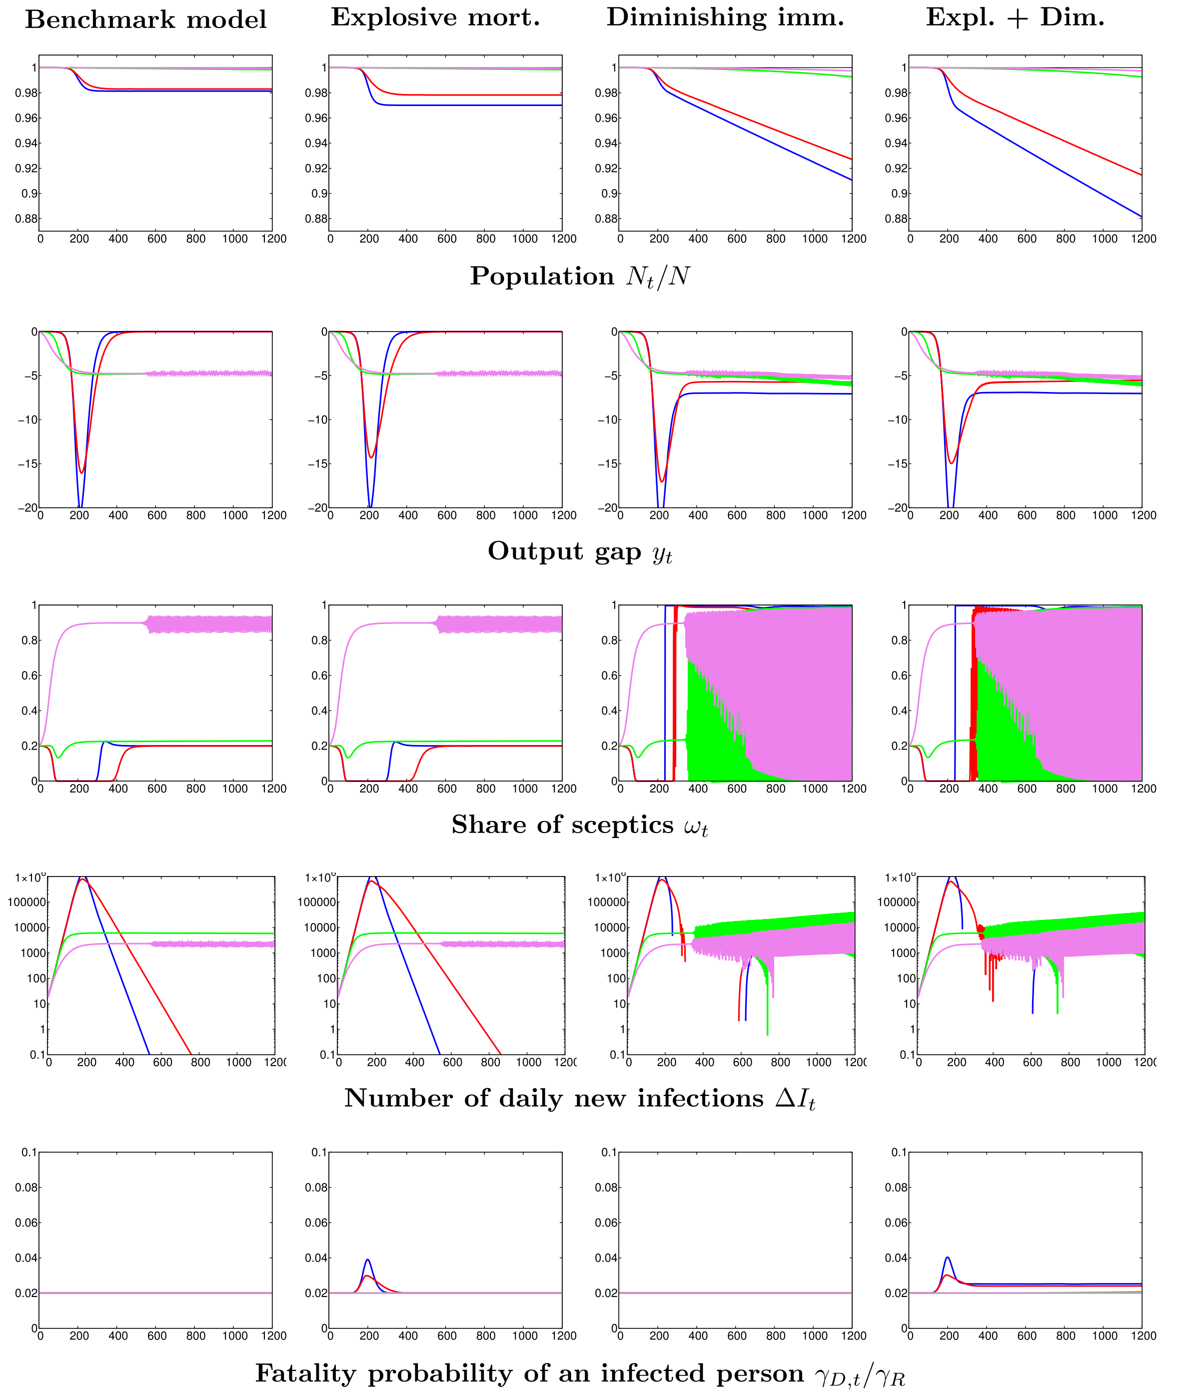

Fig. 9. Sample simulations for the four variants of the Y/N model: benchmark (left-most panels), with explosive mortality (second-left panels), diminishing immunity (second-right panels), and both explosive mortality and diminishing immunity (right-most panels). Each panel depicts four levels of government policy: no ($\phi_g = 0$; blue lines), weak ($\phi_g = 0.5$; red lines), medium ($\phi_g = 2.5$; green lines) and strong ($\phi_g = 50$; violet lines).

patients” over the ICU capacity, and in a hypothetical scenario, in which the whole population is sick, the overall death chance of a patient jumps to 10%.

Fig. 9 reports results for the four variants of the model, for four levels of the government’s response as in the previous section. It is immediately clear that the explosive mortality rate has a modest impact on its own, unless the government decides not to intervene with low values of ϕ_g. In such a case, the virus has a visibly higher death toll. Indeed, for no intervention at all with $\phi_g = 0$, the accumulated death rate jumps from 2% to 3% and around period $t = 200$, an infected person has twice as large a risk of dying

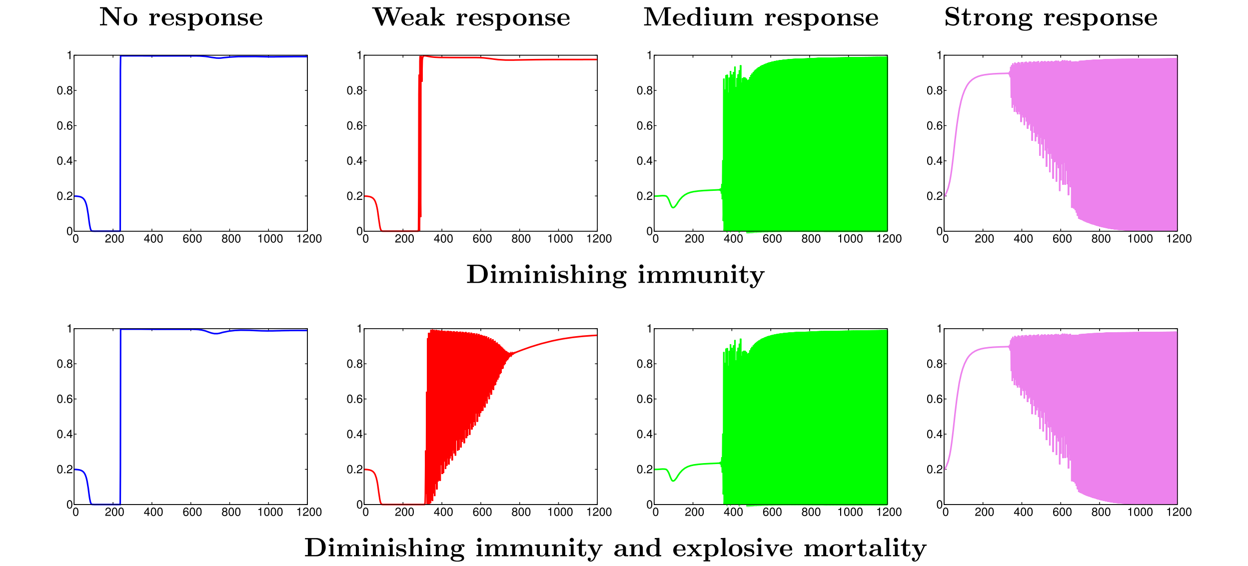

Fig. 10. Time paths of scepticism level ω_t for the model with diminishing immunity (top panels) and diminishing immunity plus explosive mortality (bottom panels), for four levels of the government policy reaction: no ($\phi_g = 0$; left-most panels), weak ($\phi_g = 0.5$; second-left panels), medium ($\phi_g = 2.5$; second-right panels) and strong ($\phi_g = 50$; right-most panels).

compared to periods in which the hospital capacity is not exceeded. Interestingly, this flash pandemic causes a less severe recession – it seems that in the peak of events, a high mortality rate eliminates a larger number of infected who would otherwise linger for longer (and many of whom would indeed survive).

A diminishing mortality rate has, on the other hand, much more severe effects on the model dynamics. Firstly, the simple trade-off between a prolonged recession and a high death toll disappears. In fact, all policies lead to a prolonged recession, which is even more severe under more lenient policies (reversing the order from the fixed immunity case). The reason for this outcome is quite intuitive. In the model with the permanent immunity of the recovered population and under a weak government response, the virus quickly burns through the population and disappears. The same response in the model with diminishing immunity means, however, that the virus burns through the population, but at some moment the population of recovered individuals becomes so large, that they can sustain a large number of people losing their immunity, which in turns yields a constant trickle of new infections and a constant hit to the economy.

The trickle of infected from recovered also implies a relatively higher accumulated death toll, for instance under the medium policy, by period $t = 1200$ a visible 0.73% of the population (represent by the green line in Fig. 9) has perished. Only a relatively strong policy could counter that, however, as seen in Fig. 10, this also leads to a potentially large share of sceptics in the population. Furthermore, strong policy results in much more volatile dynamics, as observed in Fig. 10. Interestingly, an explosive mortality rate again only exaggerates these outcomes, as the overall mortality rate stabilizes on a slightly larger level than under the fixed mortality variant of the model.

The results of diminishing mortality point to a political danger beyond the scope of this paper. It seems that in the variant of the model with diminishing immunity, there is no clear-cut trade-off between the death toll of the pandemic, and the magnitude of the recession caused by measures, which the government may want to impose to counter the pandemic. Any response leads to a recession and the constant presence of the virus, and now, it seems, the government can choose only between an enormous accumulated death toll and an unsettled political situation with a highly volatile share of sceptics. It is impossible to predict if either scenario can be sustainable, or would rather lead to massive political backlash.

6. A first empirical estimation attempt

This section outlines how the model could be empirically estimated and describes the results of a partial estimation of the core model parameters. More specifically, we first run a Monte Carlo simulation study in line with Kukacka and Sacht (2023). The authors propose a multivariate simulated maximum likelihood (SML) approach for macroeconomic Heuristic Switching Models whose estimation is infeasible under standard econometric approaches. Along with determining a set of behavioral parameters of the US economy, they also demonstrate how the SML method can reliably identify the intensity of choice parameters governing the models' nonlinear dynamics. As our behavioral SEIRD macroeconomic model shares the same belief-driven switching principle, the SML estimation method is a natural starting point for its potential empirical application. Next, we discuss the specifics of a suitable dataset that combines the US Weekly Economic Index (WEI) data and the epidemic data on new cases of COVID-19 disease. Finally, we estimate the two intensity of choice parameters from the core discrete choice switching equation (10), determining the evolution of the population share of sceptics in time.

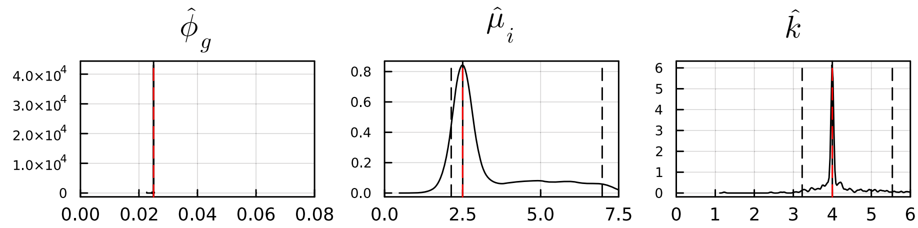

Fig. 11. Densities of the parameter estimates: the bold black curves depict the kernel density estimates of the sample densities based on 300 random runs, and the bold black vertical lines show the median point estimates. The dashed red vertical lines show the pseudo-true values, and the dashed black vertical lines depict the 95% confidence intervals of the sample estimates.

6.1. Monte Carlo study

We implement the SML estimation method by Kukacka and Sacht (2023) to the baseline model defined by equations (1) to (10) and calibrated according to Table 1. Before the empirical analysis, we first numerically investigate the capability of the SML estimator for the model in a controlled Monte Carlo simulation study. We primarily focus on estimating parameters of the core behavioral equation (7) that endogenizes the transmission rate β_t based on social distancing measures enforced by the government and the share of sceptics undermining these measures.

A few customizations of the model structure and calibration must be introduced to adapt the model for the econometric analysis. First, the deterministic structure is enriched with stochastic terms to allow for estimation. We thus add two i.i.d. idiosyncratic shocks drawn from standard normal distribution to the structure of observables: the output gap equation (8) and the expression for new cases, $G_t = \sigma E_{t-1}/100{,}000$, determining the government distancing measures. Second, the total population size is calibrated concerning the target implementation of the US data to $N = 330M$. Finally, while keeping the remaining model equations intact, we redefine the parameter of the intensity of choice related to switching towards scepticism, $\mu_y = k\mu_i$, and set $k = 4$. Such expression for μ_y as a k-multiple of μ_i does not alter the mathematical structure of the model, but it fosters the numerical optimization search. Also, the calibrated value for μ_y needs to be reduced. While its original calibration in Table 1 is well feasible for deterministic simulations in the previous sections, it leads to frequent numerical divergences of the model under the stochastic optimization search within a broad parameter space. The three pseudo-true parameters subject to estimation are the core behavioral parameters of the baseline model: the impact of the social distancing measures from equation (7), ϕ_g, and the two intensities of choice from equation (10), μ_i and $\mu_y = k\mu_i$, that drive the switching dynamics towards scepticism and back towards trusting the government.

The setup of the SML estimation algorithm follows the best practice by Kukacka and Sacht (2023). We conduct 300 independent Monte Carlo runs to support the statistical validity of the estimation output and analyze time series of 500 observations. The selected length is sufficient for a sound simulation-based analysis of the estimator properties for the model, and, at the same time, its order of magnitude is still realistic for potential datasets covering the COVID-19 pandemic period. 1,000 approximation points are generated from a multivariate normal distribution for the kernel density approximation of the likelihood in each time-point with the optimal bandwidth selection based on the Silverman (1986) rule of thumb. The standard BFGS algorithm with an automated step length and four random initial points is implemented within a constrained optimization search to minimize negative likelihood.

Fig. 11 summarizes the main results of the Monte Carlo simulation study and demonstrates a very good ability of the SML estimator to recover the coefficients of our empirical interest. Sample densities are depicted with the median point estimates in bold black and the 95% confidence intervals in dashed black. The pseudo-true values are represented by dashed red and almost perfectly overlap with the median estimates, suggesting the unbiasedness of the SML estimator for all three parameters. The impact parameter of the government social distancing measures, ϕ_g, and the relative fraction of the two intensities of choice, $k = \mu_y/\mu_i$, are recovered nearly perfectly, with minimal and very small sample variance, respectively. Recovering the remaining intensity of choice related to switching towards trusting the government, μ_i, seems the most challenging, as demonstrated by a positive skewness of the estimator distribution with an apparent fat right tail that markedly stretches the 95% confidence interval. Still, even for this parameter, the median-based SML estimator is unbiased, as demonstrated by the overlap of the point estimate with the dashed red pseudo-true value. All three parameters are thus well identified, and the SML estimator demonstrates additional favorable distributional properties. Our Monte Carlo simulation study in a controlled and perfectly specified environment thus provides substantial evidence that the model is well designed for empirical application to recover its core behavioral parameters.

6.2. Partial estimation of the core model parameters

Encouraged by the promising results of the simulation study, we now provide an initial attempt for an actual empirical application of the behavioral SEIRD macroeconomic model. The most suitable dataset for this purpose consists of US weekly data from January 25, 2020, to January 1, 2022. It covers almost two years of the onset and the most severe periods of the COVID-19 pandemic, during which the government restrictions were proactively enforced. Hence, it provides us with 102 observations combining the WEI data reported by the New York Fed[16] and the US epidemic data on new weekly COVID-19 cases.[17] WEI approximates quarterly GDP growth based on aggregated information from short-term macro-indicators such as industrial production, unemployment benefits paid, etc. The main advantage is its weekly frequency. From the WEI-implied GDP level, we approximate the output gap modeled by equation (8) using the standard Hodrick-Prescot filter with a weekly smoothing parameter $\lambda = 52^4 \cdot 6.25$ according to Ravn and Uhlig (2002). The new weekly cases then capture the newly infected people in absolute terms.

The application of the empirical dataset brings additional challenges. First and foremost, the dataset length barely exceeding 100 observations is generally questionable and might be limiting for identifying some estimated parameters and obtaining sound empirical results (Kukacka and Barunik, 2017). Second, the unprecedented drop in economic activity after the COVID-19 onset in 2020, followed by a 'recovery back' to relatively normal levels in 2021, leads to rather unusual dynamics of the resulting output gap time series. It predominantly decreases in the first half of the observed period while increasing back in the second half compared to typical cyclical fluctuations around zero. Last, not only is the dataset inevitably noisy, but the model is undoubtedly a simplified representation of the relationships driving the data-generating process of the pandemic society and economy. In the end, it turns out that a joint empirical estimation of all three parameters of interest is an overly complex task. The impact parameter of the government restriction, ϕ_g, is not identified and needs to be calibrated back to $\phi_g = 0.025$. Its problematic identification under a small empirical data sample is unsurprising as it is multiplied with both estimated intensities of choice through the interconnection between equations (7) and (10).

Finally, the two intensities of choice, μ_i and μ_y, are empirically estimated. The SML estimation procedure reveals median point estimates and the 95% sample confidence intervals $\hat{\mu}_i = 6.65\ (2.56 - 9.78)$ and $\hat{\mu}_y = 40.51\ (18.17 - 64.94)$. Hence, while the confidence bands are relatively wide compared to simulation-based results, both intensities of choice are statistically significant at the 5% level. Importantly, they are also of the same orders of magnitude as the calibrated values of μ_i and μ_y utilized for simulations in the previous sections (cf. equation (7)). The empirical results thus provide statistical evidence in favor of evolutionary behavioral switching towards scepticism and back towards trusting the government during the US pandemic of 2020 and 2021.

7. Concluding remarks

In this study, we investigated the role of sceptics in the effectiveness of governmental containment policies through a reduced-form specification of the trade-off between the public health and the economic dimensions in an epidemic such as COVID-19. Despite the parsimony of our behavioral SEIRD model, we could model in a reasonable manner how policy or disease scepticism affects the transmission rate, the reproduction rate, and thus the general evolution of the disease, as well as the economic sphere. In our framework, the level of scepticism is endogenously determined by the state of the pandemic and the economy, which leads to four interesting findings.

Firstly, we show that in the presence of sceptics, the government can essentially choose one out of two scenarios: a flash pandemic, in which everyone gets infected (with the corresponding high death toll), but the economy quickly recovers; or containing the virus, but at the cost of a prolonged economic downturn and a constant trickle of new infections. Ironically, the latter choice fuels scepticism, since strict lockdown measures prevent the spread of the disease and cause a recession. Secondly, the specific magnitude of the policy reaction plays a little role, since the virus spreads quickly enough to force an even more cautious government to eventually impose a strict lockdown. These two results showed the particular viciousness of the virus in our model, which was calibrated to the initial Sars-Cov-2 variants.

Thirdly, it is more difficult to impose the lockdown when the government and the general public are misaligned in terms of their policy goals. In particular, when the government reacts to the incidence rate, whereas the public cares more about the deaths (which lag after the incidence rate) and the output gap (which is inevitably brought by the lockdown), it is visibly more challenging to manage the sceptics. This result indicates that real-world lockdown measures need to be backed by clear communication between the government and public about the state of the pandemic.

Next, our simulations suggest that ICU capacity constraints play little role as long as the government is willing to impose a strict lockdown. However, if the virus immunity can wane over time (as is the case of COVID-19), loose government policy becomes less effective than the harsh one. When the government keeps the economy open, recovered individuals gradually return to the pool of susceptible population and become infected again, which puts the economy in a state of permanent recession due to labor shortages. In our calibration, this recession is, in fact, harsher than the one caused by restrictive lockdowns, and, in addition, it is accompanied by a harrowing death toll. Overall, this result suggests that the "just let it go" policy can be both dangerous and counter-productive.

Finally, we outline and analyze how the model could be empirically estimated via the multivariate simulated maximum likelihood method. Using a dataset combining the US WEI data and epidemic data on new cases of COVID-19 disease, we subsequently estimate the two intensity of choice parameters driving behavioral switching towards scepticism and back towards trusting the government. Both coefficients are statistically significant at the 5% level and have the same orders of magnitude as the calibrated values utilized for simulations. These empirical results provide sound statistical evidence in favor of the evolutionary dynamics of the population share of sceptics during the US pandemic of 2020 and 2021.

The current framework can be extended in various directions. First and foremost, it would be interesting re-calibrate the model on the basis of real empirical data. Furthermore, one could introduce the public health vs. economic trade-off in the government's

[16] Available at www.newyorkfed.org/research/policy/weekly-economic-index.

[17] Available at ourworldindata.org/covid-cases.

policy function, as well as model the economic dimension in a more detailed manner. Finally, a natural extension of the model is to include vaccinations, which could be subject to endogenously formed scepticism on their own. We intend to tackle these and other issues in future work.

Declaration of competing interest

There is no competing interest.

Data availability

Data will be made available on request.

Appendix A. Intensity of choice

Warning Figs. A.1–A.4 in this Appendix represent simulations in which the intensity of choice for the pull force was rescaled such that $\mu_i = 1000\hat{\mu}_i$. This rescalation was done in order to secure the numerical stability of the simulations.

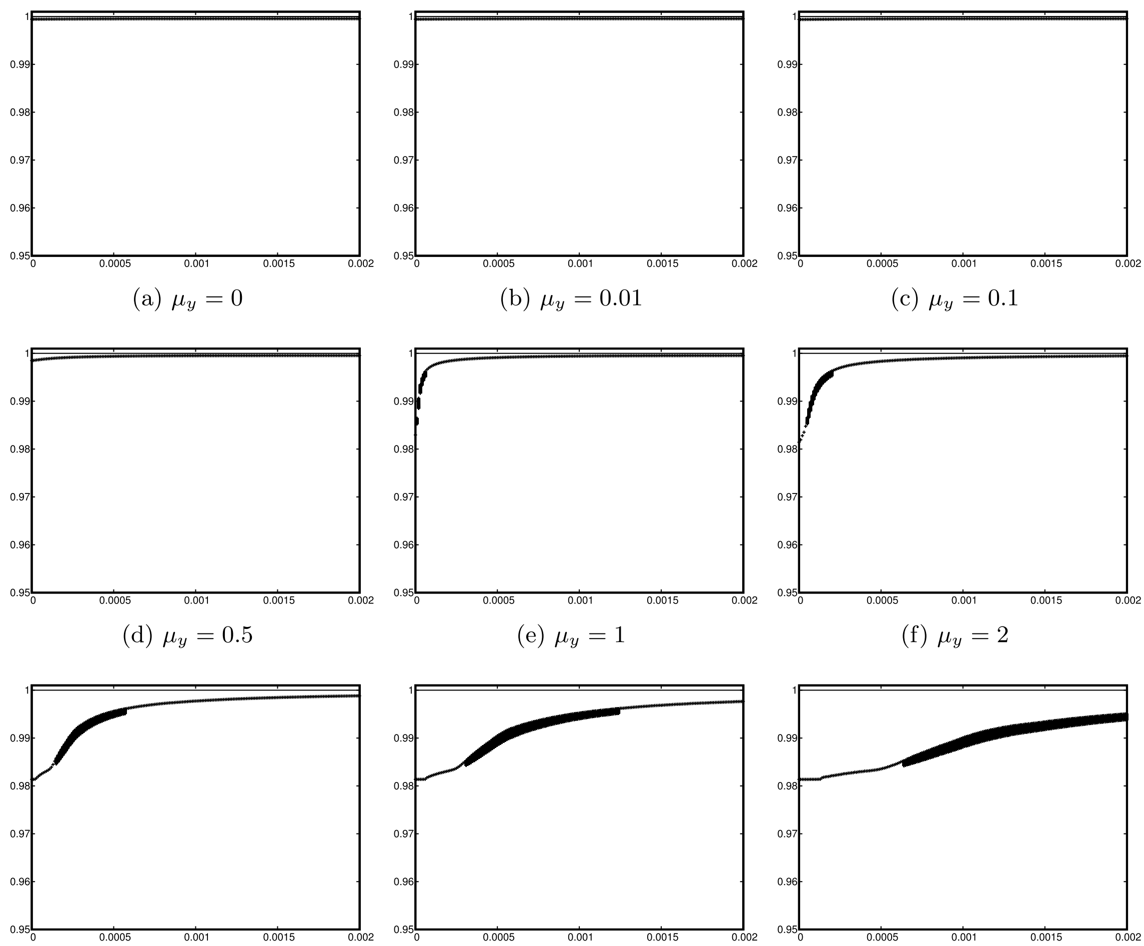

(a) $\mu_y = 0$ (b) $\mu_y = 0.01$ (c) $\mu_y = 0.1$

(d) $\mu_y = 0.5$ (e) $\mu_y = 1$ (f) $\mu_y = 2$

(g) $\mu_y = 5$ (h) $\mu_y = 10$ (i) $\mu_y = 20$

Fig. A.1. Bifurcation diagram. Population as a function of $\hat{\mu}_i$, for a given value of μ_y. Each vertical slice represents the range of outcomes in periods $t \in \{1101, \ldots, 1200\}$ for the given parameter $\hat{\mu}_i$.

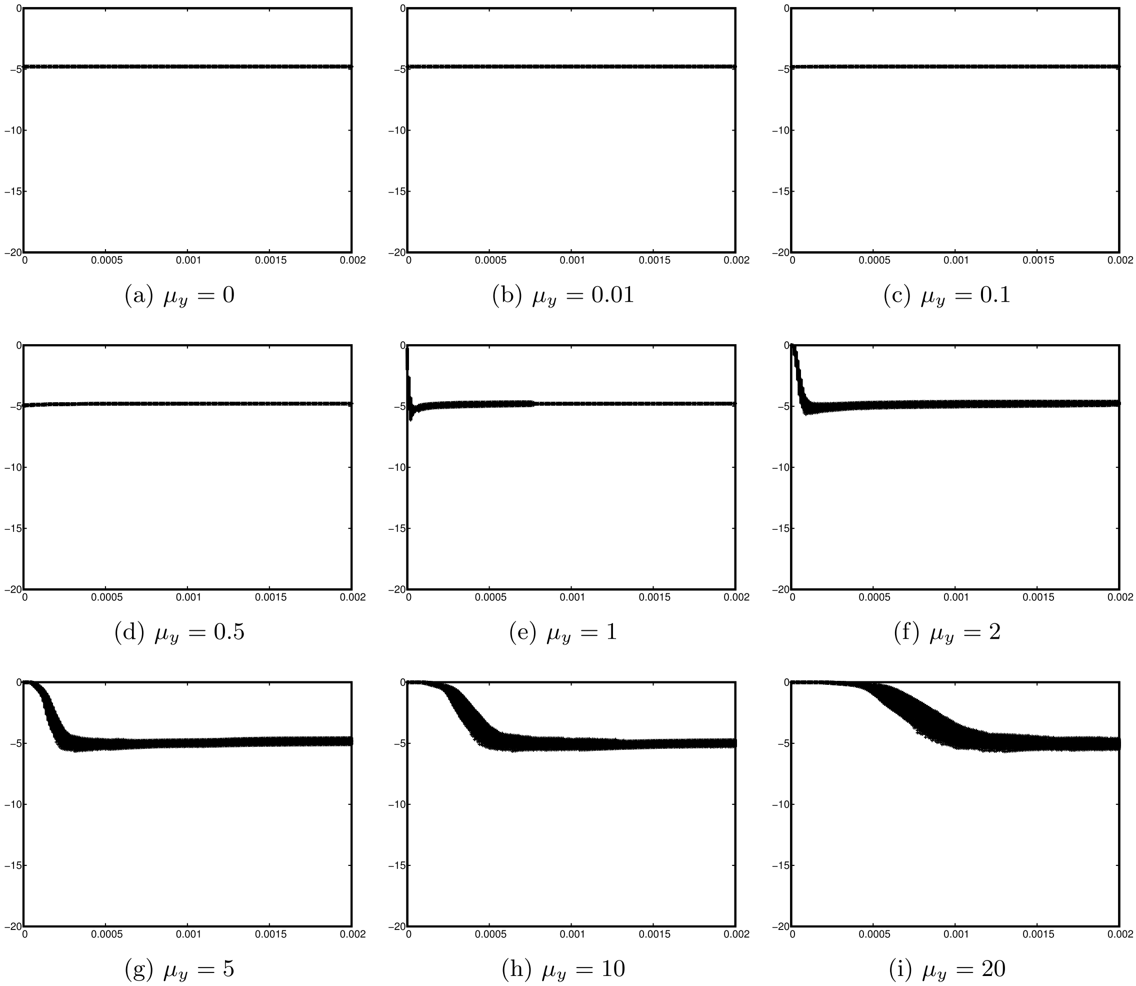

Fig. A.2. Bifurcation diagram. Output gap as a function of $\hat{\mu}_i$, for a given value of μ_y. Each vertical slice represents the range of outcomes in periods $t \in \{1101, \ldots, 1200\}$ for the given parameter $\hat{\mu}_i$.

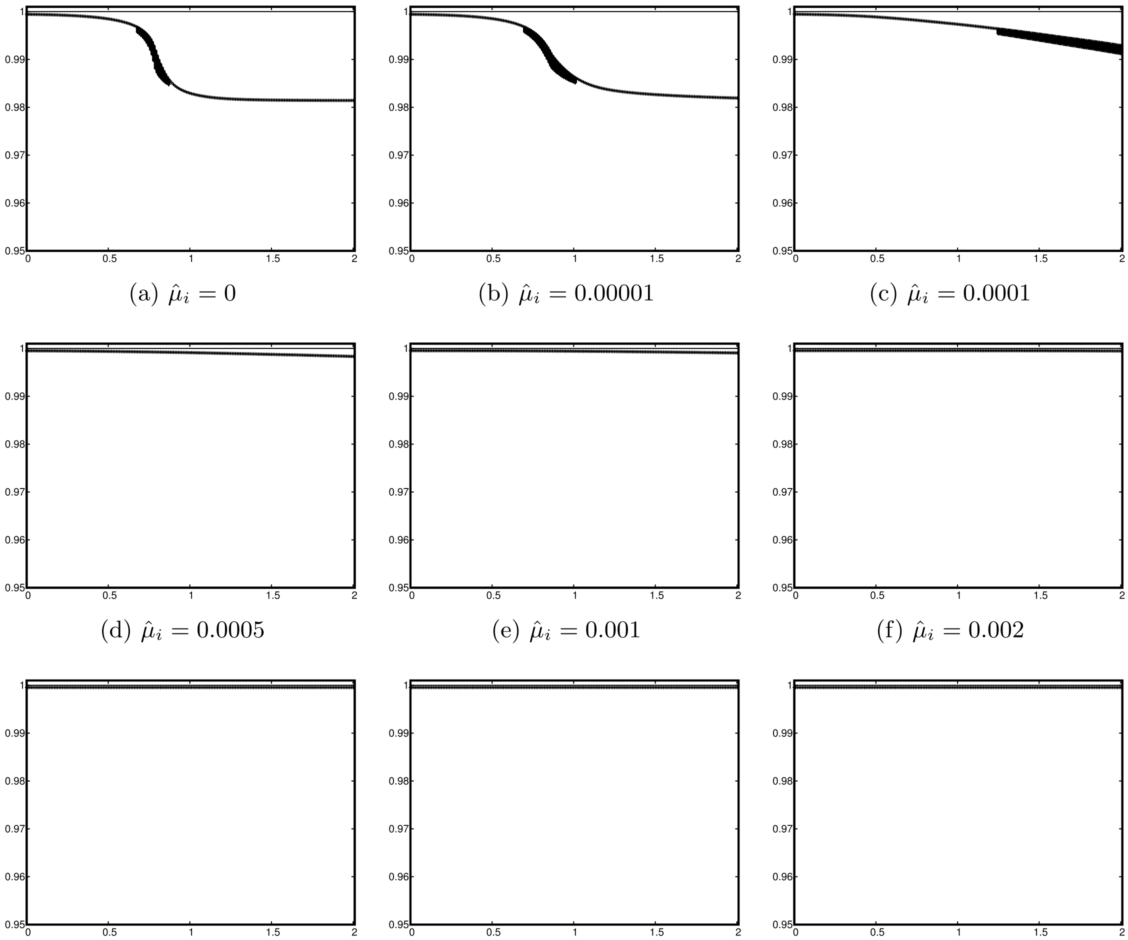

(a) $\hat{\mu}_i = 0$ (b) $\hat{\mu}_i = 0.00001$ (c) $\hat{\mu}_i = 0.0001$

(d) $\hat{\mu}_i = 0.0005$ (e) $\hat{\mu}_i = 0.001$ (f) $\hat{\mu}_i = 0.002$

(g) $\hat{\mu}_i = 0.005$ (h) $\hat{\mu}_i = 0.01$ (i) $\hat{\mu}_i = 0.02$

Fig. A.3. Bifurcation diagram. Population as a function of μ_y, for a given value of $\hat{\mu}_i$. Each vertical slice represents the range of outcomes in periods $t \in \{1101, \dots, 1200\}$ for the given parameter μ_y.

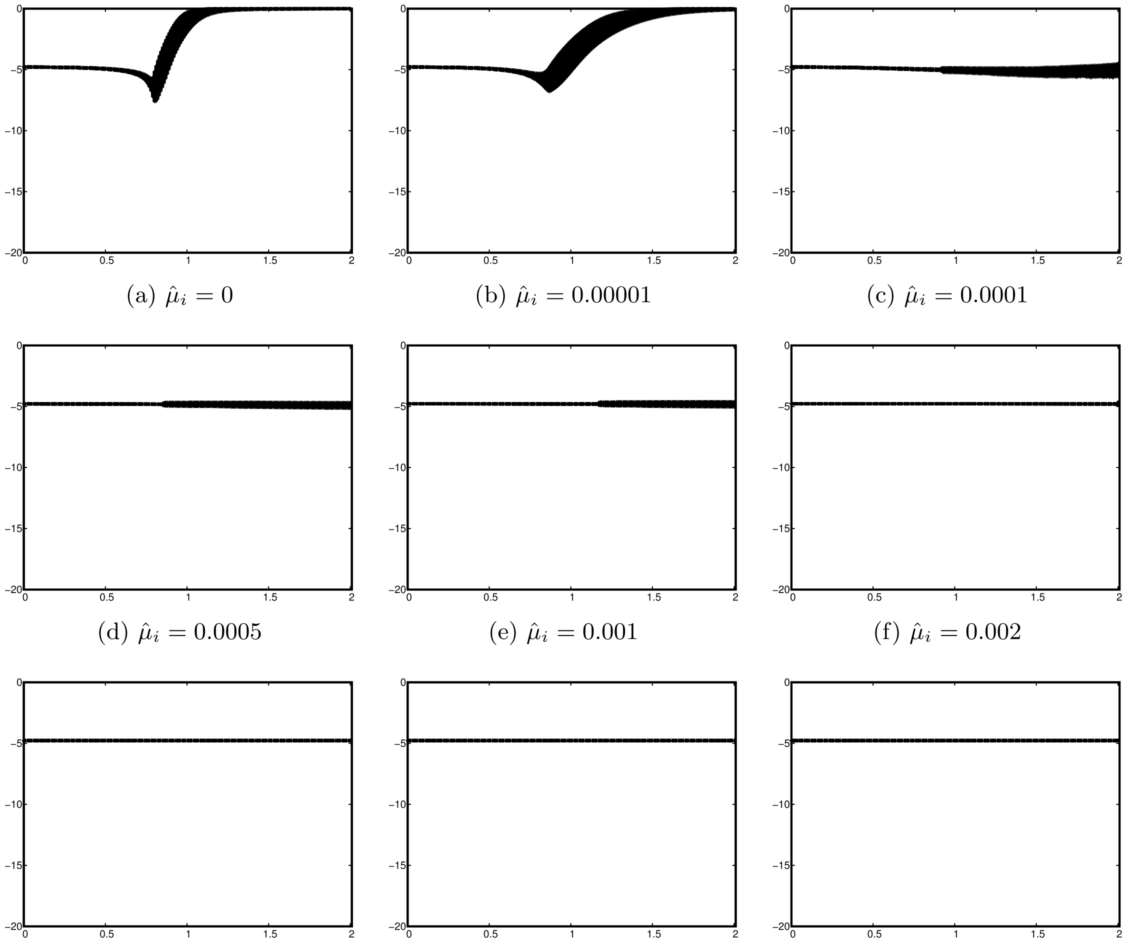

(a) $\hat{\mu}_i = 0$ (b) $\hat{\mu}_i = 0.00001$ (c) $\hat{\mu}_i = 0.0001$

(d) $\hat{\mu}_i = 0.0005$ (e) $\hat{\mu}_i = 0.001$ (f) $\hat{\mu}_i = 0.002$

(g) $\hat{\mu}_i = 0.005$ (h) $\hat{\mu}_i = 0.01$ (i) $\hat{\mu}_i = 0.02$

Fig. A.4. Bifurcation diagram. Output gap as a function of μ_y, for a given value of $\hat{\mu}_i$. Each vertical slice represents the range of outcomes in periods $t \in \{1101, \dots, 1200\}$ for the given parameter μ_y.

References

Aguiar, M., Amador, M., Farhi, E., Gopinath, G., 2013. Crisis and commitment: Inflation credibility and the vulnerability to sovereign debt crises. Working Paper 19516. National Bureau of Economic Research.
Allcott, H., Boxell, L., Conway, J., Gentzkow, M., Thaler, M., Yang, D., 2020. Polarization and public health: partisan differences in social distancing during the coronavirus pandemic. J. Public Econ. 191, 104254.
Assenza, T., Heemeijer, P., Hommes, C.H., Massaro, D., 2021. Managing self-organization of expectations through monetary policy: a macro experiment. J. Monet. Econ. 117, 170–186.
Atkeson, A., 2020. What will be the economic impact of COVID-19 in the US? rough estimates of disease scenarios. NBER Working Paper 26867. National Bureau of Economic Research.
Atkeson, A.G., Kopecky, K., Zha, T., 2021. Behavior and the transmission of COVID-19. Am. Econ. Rev. Pap. Proc. 111, 356–360.
Barro, R.J., Gordon, D.B., 1983. Rules, discretion and reputation in a model of monetary policy. J. Monet. Econ. 12 (1), 101–121.
Basurto, A., Dawid, H., Harting, P., Hepp, J., Kohlweyer, D., 2020. Economic and epidemic implications of virus containment policies: insights from agent-based simulations. Bielefeld Working Papers in Economics and Management 05-2020. Bielefeld University.
Basurto, A., Dawid, H., Harting, P., Hepp, J., Kohlweyer, D., 2023. How to design virus containment policies? A joint analysis of economic and epidemic dynamics under the COVID-19 pandemic. J. Econ. Interact. Coord. 18 (2), 311–370.
Brock, W.A., Hommes, C.H., 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. J. Econ. Dyn. Control 22 (8–9), 1235–1274.
Bursztyn, L., Rao, A., Roth, C., Yanagizawa-Drott, D.H., 2020. Misinformation during a pandemic. NBER Working Paper 27417. National Bureau of Economic Research.
Charron, N., Lapuente, V., Rodríguez-Pose, A., 2023. Uncooperative society, uncooperative politics or both? Trust, polarization and populism explain excess mortality for COVID-19 across European regions. Eur. J. Polit. Res. 62 (3), 781–805.
Cliffe, M., 2020. Pandenomics. How COVID-19 could change the world. Think economic & financial analysis. ING. https://think.ing.com/reports/pandenomics-how-covid-19-could-change-the-world.
Dasaratha, K., 2020. Virus dynamics with behavioral responses. arXiv:2004.14533. https://ideas.repec.org/p/arx/papers/2004.14533.html, 2020.
Delli Gatti, D., Reissl, S., 2020. ABC: an agent based exploration of the macroeconomic effects of Covid-19. Working Paper 8763. CESifo.
Di Guilmi, C., Galanis, G., Basokos, G., 2022. A behavioural SIR model and its implications for physical distancing. Rev. Behav. Econ. 9 (1), 45–63.
Eichenbaum, M., Rebelo, S., Trabandt, M., 2021. The macroeconomics of epidemics. Rev. Financ. Stud. 34 (11), 5149–5187.
Eksin, C., Paarporn, K., Weitz, J., 2019. Systematic biases in disease forecasting. Epidemics 27, 96–105.

Evans, G.W., Honkapohja, S., 2012. Learning and Expectations in Macroeconomics. Princeton University Press, Princeton, NJ.
Faia, E., Fuster, A., Pezone, V., Zafar, B., 2022. Biases in information selection and processing: survey evidence from the pandemic. Rev. Econ. Statist., 1–46.
Fauci, A., Lane, C., Redfield, R., 2020. Covid-19 – navigating the uncharted. N. Engl. J. Med. 382 (13), 1268–1269.
Flaschel, P., Veneziani, R., Tavani, D., Galanis, G., 2022. Pandemics and aggregate demand: a framework for policy analysis. Rev. Behav. Econ. 9 (1), 1–44.
Funk, S., Bansal, S., Bauch, C., Eames, K., Edmunds, W., Galvani, A., Klepac, P., 2015. Nine challenges in incorporating the dynamics of behaviour in infectious diseases models. Epidemics 10, 21–25.
Hommes, C., 2013. Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems. Cambridge University Press, Cambridge, UK.
Hommes, C., Lustenhouwer, J., 2019a. Inflation targeting and liquidity traps under endogenous credibility. J. Monet. Econ. 107, 48–62.
Hommes, C., Lustenhouwer, J., 2019b. Managing unanchored, heterogeneous expectations and liquidity traps. J. Econ. Dyn. Control 101, 1–16.
Kaplan, G., Moll, B., Violante, G.L., 2020. The great lockdown and the big stimulus: Tracing the pandemic possibility frontier for the US. NBER Working Paper 27794. National Bureau of Economic Research.
Kermack, W.O., McKendrick, A.G., 1927. A contribution to the mathematical theory of epidemics. Proc. R. Soc. Lond. Ser. A, Contain. Pap. Math. Phys. Character 115 (772), 700–721.
Kukacka, J., Barunik, J., 2017. Estimation of financial agent-based models with simulated maximum likelihood. J. Econ. Dyn. Control 85, 21–45.
Kukacka, J., Sacht, S., 2023. Estimation of heuristic switching in behavioral macroeconomic models. J. Econ. Dyn. Control 146, 104585.
Lohmann, S., 1992. Optimal commitment in monetary policy: credibility versus flexibility. Am. Econ. Rev. 82 (1), 273–286.
Lux, T., 2021. The social dynamics of Covid-19. Phys. A, Stat. Mech. Appl. 567, 125710.
Massaro, D., 2013. Heterogeneous expectations in monetary DSGE models. J. Econ. Dyn. Control 37 (3), 680–692.
Matějka, F., McKay, A., 2015. Rational inattention to discrete choices: a new foundation for the multinomial logit model. Am. Econ. Rev. 105 (1), 272–298.
Mellacher, P., 2023. The impact of corona populism: empirical evidence from Austria and theory. J. Econ. Behav. Organ. Forthcoming.
Milosh, M., Painter, M., Sonin, K., Van Dijcke, D., Wright, A.L., 2021. Unmasking partisanship: polarization undermines public response to collective risk. J. Public Econ. 204, 104538.
Mokhtarzadeh, F., Petersen, L., 2021. Coordinating expectations through central bank projections. Exp. Econ. 24 (3), 883–918.
Msemburi, W., Karlinsky, A., Knutson, V., Aleshin-Guendel, S., Chatterji, S., Wakefield, J., 2022. The WHO estimates of excess mortality associated with the COVID-19 pandemic. Nature 613, 130–137.
Ravn, M.O., Uhlig, H., 2002. On adjusting the Hodrick-Prescott filter for the frequency of observations. Rev. Econ. Stat. 84 (2), 371–376.
Rossman, H., Meir, T., Somer, J., Shilo, S., Gutman, R., Arie, A.B., Segal, E., Shalit, U., Gorfine, M., 2021. Hospital load and increased COVID-19 related mortality in Israel. Nat. Commun. 12 (1), 1–7.
Silverman, B.W., 1986. Density Estimation for Statistics and Data Analysis. Chapman and Hall, London.
Svensson, L.E., 2020. Monetary policy strategies for the Federal Reserve. Int. J. Cent. Bank. 16, 133–193.
Townsend, J.P., Hassler, H.B., Wang, Z., Miura, S., Singh, J., Kumar, S., Ruddle, N.H., Galvani, A.P., Dornburg, A., 2021. The durability of immunity against reinfection by Sars-Cov-2: a comparative evolutionary study. Lancet Microbe 2 (12), e666–e675.
Wang, H., Wang, Z., Dong, Y., Chang, R., Xu, C., Yu, X., Zhang, S., Tsamlag, L., Shang, M., Huang, J., Wang, Y., Xu, G., Shen, T., Zhang, X., Cai, Y., 2020. Phase-adjusted estimation of the number of coronavirus disease 2019 cases in Wuhan, China. Cell Discov. 6 (10), 1–8.