DIW BERLIN

1861

Discussion Papers

Deutsches Institut für Wirtschaftsforschung 2020

Viral Shocks to the World Economy

Konstantin A. Kholodilin and Malte Rieth

IMPRESSUM

DIW Berlin
German Institute for Economic Research
Mohrenstr. 58
10117 Berlin

Tel. +49 (30) 897 89-0
Fax +49 (30) 897 89-200
https://www.diw.de

ISSN electronic edition 1619-4535

Papers can be downloaded free of charge from the DIW Berlin website:
https://www.diw.de/discussionpapers

Discussion Papers of DIW Berlin are indexed in RePEc and SSRN:
https://ideas.repec.org/s/diw/diwwpp.html
https://www.ssrn.com/link/DIW-Berlin-German-Inst-Econ-Res.html

Viral shocks to the world economy

Konstantin A. Kholodilin[a,b], Malte Rieth[b,c]

[a]*DIW Berlin, Berlin, Germany.*
[b]*National Research University Higher School of Economics, St. Petersburg, Russia.*
[c]*Corresponding author: Martin-Luther-Universität Halle-Wittenberg, Universitätsring 3, 06108 Halle, Germany; mrieth@diw.de.*

First version: April 7, 2020 *This version: December 8, 2021* *Conflict of interest: none.*

Abstract

We construct news-based viral disease indices and study the economic consequences of epidemics on the world economy, using structural vector autoregressions. Epidemic shocks have significantly and persistently negative effects on output and prices that last for 2-3 quarters. There is no quick recovery and no overshooting. The output losses are permanent. Country studies show that the direct effects are four times larger than the indirect effects through trade and financial markets. Furthermore, the simultaneous fall of output and prices suggests that demand-side dominate supply-side contractions and that fiscal stimulus is an appropriate crisis response to minimize permanent output losses.

Keywords: Epidemics, text analysis, fiscal policy, structural vector autoregressions, Covid-19, world economy.

JEL codes: C32, E62, H12, I18.

1. Introduction

The worldwide spreading of the coronavirus in 2020/21 is a large risk for human lives and the world economy. It causes severe disruptions in both the production and the service sector due to mortality and morbidity, quarantines, lockdowns, travel restrictions, as well as changes in investor and consumer behavior. Policy makers face difficulties in understanding the pandemic and finding appropriate responses. How long will the downturn last? Will there be an overshooting of production, making up for the losses during the epidemic, or should fiscal policy step-in immediately? Are epidemics more demand shocks or supply shocks? Answering these questions now can help public policy in tailoring the size, duration, and type (demand-side vs. supply-side) of interventions as an analysis of the economic dynamics following the Covid-19 pandemic can only be conducted many years after it occurred, which does not inform the urgent debate.

In this paper, we construct monthly news-based viral disease indices and study the effects of epidemic shocks on the world economy, using structural vector autoregressions. We analyze the text of over 500 million documents by counting words like 'SARS', 'swine flu', or 'influenza', following the approach of Baker et al. (2016) for measuring economic policy uncertainty. Thereby, we address the problem that monthly (or quarterly) frequency data on infected or fatalities of past epidemics are not consistently available over time and across countries. In contrast, our indices are available at the monthly frequency for between 30-100 years, depending on the media considered. In an external validation of the indices, we show that they track the timing of epidemics as defined by the World Health Organization (WHO) and provide a close mapping to selectively available clinical data. Moreover, the indices are useful for analyzing the economic impact of pulmonary diseases as they also reflect the public awareness, which is crucial for private and public decisions. We aim at providing the indices to other interested scholars.

We use the viral disease indicators to trace out the global economic impact of epidemic shocks. First, we estimate and compare the worldwide output effects of two major virus outbreaks in the 21st century, the severe acute respiratory syndrome coronavirus (SARS-CoV) outbreak in East Asia in 2002/03 and the swine flu pandemic (influenza A/H1N1) of North America in 2009/10. The findings suggest that both shocks reduce economic activity significantly in the country of origin, in many other countries, and globally for two quarters or more. We contrast these findings with the dynamic impact of seasonal influenza shocks to assess a popular contention that Covid-19 has the same impact as a flu. The estimated effects of regular flu, while statistically significant, are only short-lived and much smaller.

Then, we construct an aggregate news-based index for influenza-like diseases and assess the international repercussions of epidemic shocks. We find that world output falls immediately and that it remains below trend for several quarters. It recovers only gradually and there is no evidence of overshooting, suggesting that epidemics entail permanent output losses. World trade also declines significantly for about six months, mostly driven by lower imports and exports of advanced economies. World equity prices and employment fall persistently as well. Despite a temporary increase in world retail sales, consumer prices decline globally. In an extensive sensitivity analysis, we show that the results are robust to measurement error, reverse causality, omitted variables, alternative index construction, estimation methodology, and control variables. To gauge the direct and indirect effects of the shocks, we conduct two country studies. We compare Hong Kong, which was most affected by SARS, and Germany, a small open economy that was only indirectly hit by epidemics before Covid-19. Qualitatively, the responses of both economies to the shocks are similar, mirroring the dynamics of world aggregates. However, the economic damage for Hong Kong is about four times larger than for Germany.

The baseline sample is 1991M1-2019M12, excluding the Covid-19 pandemic. Otherwise the associated extreme outliers would dominate all other information in the sample. The structural VAR would simply match the data in 2020. To see how well the model and results are transferable to Covid-19 and historical pandemics, we perform two out-of-sample validations. First, we confront the model's predictions with the actual data in 2020. We show that the estimated impulse responses to sensibly calibrated epidemic shocks match the observed production losses due to Covid-19 well, both qualitatively and quantitatively. Second, we conduct a historical analysis of pandemic shocks in the United States since 1920 using an alternative media sample and index. This time span adds the influenza outbreak of 1929, the Asian flu of 1957/58, and the Hong Kong flu of 1968/69 to the sample, among other smaller influenza outbreaks. These findings also confirm the evidence for the world economy based on the baseline sample. Production and prices drop significantly in response to an adverse aggregate health shock.

In public economics, there is a tradition of work on the impact of viral diseases. Focusing on the 1918 Spanish flu, Almond (2006) assesses long-term cohort outcomes in the U.S., Karlsson et al. (2014) document economic performance in Sweden, and Correia, Luck and Verner (2020) identify the economic and policy responses of U.S. areas. Adda (2016) investigates the interaction between viral diseases and economic activity in France. There is also a rapidly growing literature on the implications of Covid-19 for public policy (Marinescu, Skandalis and Zhao, 2021; Mitman and Rabinovich, 2021; Clemens and Veuger, 2021). These studies provide precisely identified effects using microeconomic methods and focusing on certain diseases or single countries. By contrast, our interest are the average global effects of epidemics and we use multiple time-series analysis.

In general equilibrium analysis, the literature on epidemics was thin before Covid-

19 and largely confined to model-based approaches. McKibbin, Sidorenko et al. (2006), Dixon, Lee, Muehlenbeck, Rimmer, Rose and Verikios (2010), and Verikios, Sullivan, Stojanovski, Giesecke and Woo (2016) use computable general equilibrium models for quarterly or annual frequencies. While these models facilitate a detailed analysis of different sectors and countries, they are originally designed to study comparative-statics. For a short-run dynamic analysis of epidemics in individual countries, Keogh-Brown, Wren-Lewis, Edmunds, Beutels and Smith (2010) employ a semi-structural model for the U.K. and Eichenbaum, Rebelo and Trabandt (2020) develop a dynamic-stochastic general equilibrium model for the U.S. An advantage of the model-based approaches over empirical work is that they allow for theoretical insights into policy interventions (Bisin and Gottardi, 2021). Limitations are that they require behavioral and parametric assumptions as well as calibrated shocks to mimic the intensity of epidemics. Empirical approaches are adopted by Barro, Ursúa and Weng (2020), who estimate the global impact of the Spanish flu using annual death rates, and Jordà, Singh and Taylor (2020), who focus on the longer-term economic consequences of pandemics based on annual data.

We complement the insights from these studies by estimating the global economic consequences of epidemics using text analysis and monthly data. To the best of our knowledge, this is the first paper that constructs news-based indices for viral diseases. We provide the indices to be used by other researchers in the future and exploit them ourselves to document three novel stylized facts with direct policy implications. First, the economic damage of epidemics is long-lasting. There is no quick recovery and no overshooting, belying hopes of a swift return to normal capacity after Covid-19 and suggesting that the output losses are permanent. Second, the economic disruptions in countries directly affected by an epidemic are about four times larger than the harm to countries indirectly affected via supply chains, trade, financial markets, or confidence. This finding underscores the importance of early health interventions that limit international contagion. Third, the negative demand effects of epidemics dominate the adverse supply effects: economic activity and prices drop simultaneously. This pattern indicates that fiscal stimulus is an appropriate reaction to epidemic outbreaks in order to curb the otherwise permanent output losses.

2. Data, construction and validation of viral disease indices

In this section, we first explain the derivation of the news-based viral disease indices (Section 2.1). Then, we discuss their benefits and drawbacks relative to direct clinical measures of diseases (Section 2.2). Finally, we validate the news indices using such direct measures, when these are available (Section 2.3)

2.1. Derivation of news-based viral disease indices

We measure the occurrence and intensity of viral diseases through news-based indices. To construct them, we use automated text analysis of two online media archives. The first one is the database Genios.[1] It includes about 2200 high-quality German-speaking media between 1990M1 and 2021M10 with the total number of documents exceeding 500 million. The database contains a wide variety of media, including daily press, specialized weekly journals, and magazines for general public as well as publications devoted to specific firms and persons. The second source is the archive of The New York Times (NYT), which spans the period 1851M1-2021M10.[2] We search for the following five keywords, which capture epidemics during our sample period due to pulmonary diseases: 'SARS,' 'swine flu' ('Schweinegrippe' in German), 'MERS,' 'Influenza' ('Grippe' in German), and 'Coronavirus'. Since the wording of the 2020-21 pandemic has changed over time, we alternatively search for 'Covid-19' which has gained importance. We count the monthly occurrences of these keywords.

As both databases are extremely large, we use the statistical programming language **R** to automatize the collection of the information. For each keyword, we create a sequence of daily dates which is supplied together with the keyword as an input to the database query. In the free of charge version of both archives, the automated search for the keyword and date results in a list of media items. These contain the keyword in question as well as short excerpts of each item, including typically the first few lines of it. The number of occurrences is stored and then the next date with the same keyword is supplied. In such a way, for each keyword and date we obtain the occurrences of the keyword under inspection. The automated text search needs to be tailored specifically to each archive, as their search syntax differs.

[1] www.genios.de.

[2] https://www.nytimes.com/search?.

Given that the number of texts collected in the databases changes over time, we normalize the plain counts. We divide them by the number of occurrences of the word 'der', which is the most widely used word in German language, in the case of Genios, and by the occurrence of the word 'new' for the NYT archive, since the word 'the' is not searchable in that database. Thus, the index for keyword i is computed as follows:

$$A_t^i = 1000 \times \frac{N_t^i}{N_t^{\text{word}}}$$

where N_t^i is the number of occurrences of the i-th keyword in month t and N_t^{word} is the number of occurrences of the normalizing word. The resulting value is a relative frequency of use of the keyword in the media. The normalization renders the values comparable over time and across sources.

The black lines in Figure 1 show the news-based disease indicators. These reflect both the timing of the public awareness of epidemics (when the news about it start appearing in the media) and the intensity of media coverage (the height of the series). Media coverage reflects the relevance of the disease from the standpoint of reporting media, summarizing all publicly available information about the likely severity of a disease.[3] Comparing the maxima of the SARS, swine flu, MERS, and coronavirus indices shows that the Covid-19 pandemic is 1-2 orders of magnitude larger than the others.

The deadliest diseases in the sample are avian flu and MERS. The mortality rates are 60% and 36%, respectively. SARS and Covid-19 have, at the time of writing, expected case-fatality rates of roughly 10% and <1%, respectively (Fauci et al., 2020). Swine and seasonal flus, on the other hand, are estimated to be deadly for 0.4% and 0.1% of the infected, respectively. The transmissibility also differs markedly across diseases. The estimated reproduction number, which is defined as the expected number of infected cases generated by one infected case, for SARS, Covid-19, swine flu, seasonal influenza, avian flu, and MERS is 3, 2-3, 1-2, 1, <1, and <1, respectively. Online Appendix B provides further details of the diseases.

[3] As the media mention the word 'MERS' sometimes in other contexts before the emergence of a disease with this name in 2012, we set index values before 2011 to zero. Moreover, there are small spikes for the SARS and MERS index towards the end of the sample in conjunction with the Covid-19 pandemic. Setting these values to zero as well, leaves all the results virtually unchanged.

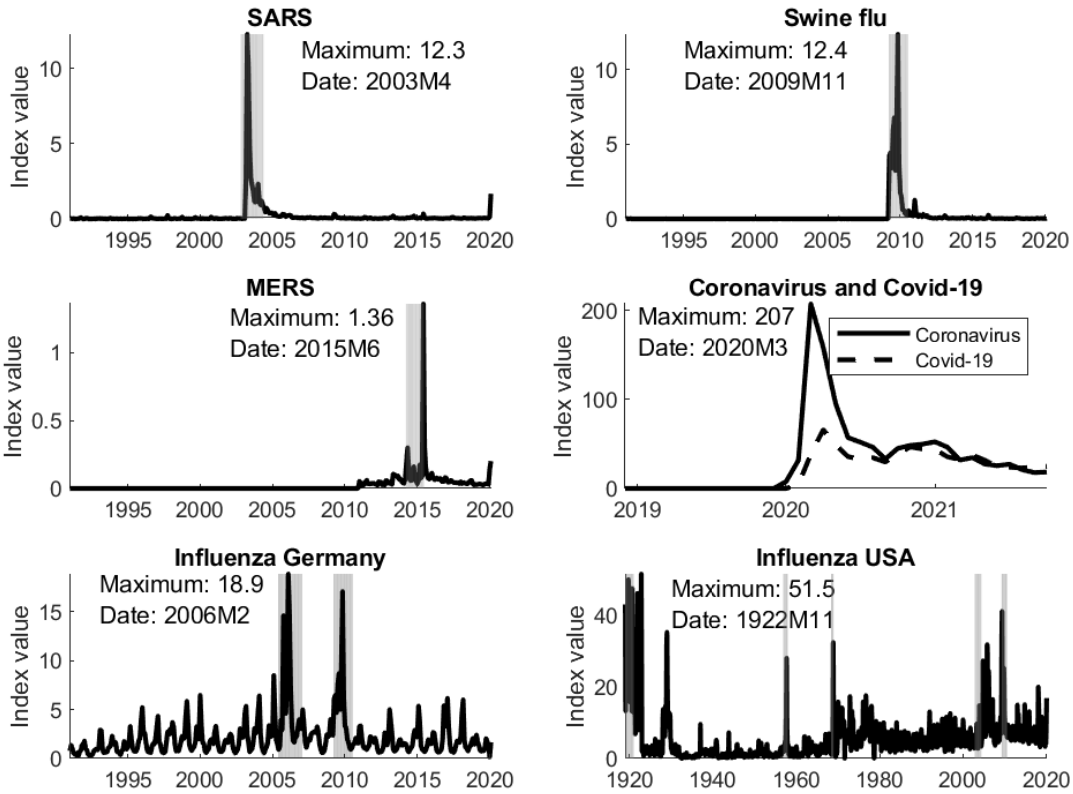

Figure 1: News-based disease indices. *Notes:* The figure shows news-based disease indicators based on text analysis of the Genios and NYT media archives. The frequency of keyword appearances are normalized. Each subplot also lists the maximum of each series and the corresponding months. The shaded areas indicate the retrospective definition of the WHO of start and ending dates of an epidemic or pandemic.

2.2. Advantages and disadvantages of news-based viral disease indices

The news-based indices of viral diseases have some advantages over clinical measures. Most importantly, they are available at a relatively high frequency, that is, monthly, and for between 30-100 years. This allows estimating time-series models on macroeconomic data. Another main argument for using news-based indices is that they capture the public awareness of contagious diseases which, in turn, determines the response of the private sector and of public authorities to the outbreaks.

In contrast, statistics on laboratory-confirmed cases are available only annually for longer samples, or they are weekly/monthly but fragmented and available for at most 10 years and a few countries. For example, the U.S. Center for Disease Control and Prevention publishes time series of reported cases only for seasonal diseases, such as flu, for the period 1999-2018. Our index for the U.S. adds 80 years to that. Data on total deaths are available for longer periods. A prominent source is the United Nations

(UNdata) that provides the monthly time series 'Deaths by month of death'. While this series is available for the period 1980-2020 for several countries, it is often incomplete and not always up-to-date. The national statistical offices sometimes provide consecutive and more timely information. That in, turn, can differ from the UN figures and covers shorter periods, typically starting around 2000. A few nice exceptions are Switzerland that publishes monthly data on total deaths starting from 1901, France from 1946, and Germany from 1950. Weekly or daily time series of total deaths are much shorter. For example, the weekly data for the UK start in 2010, while daily data for Denmark and Sweden start in 2007 and 2015, respectively. Finally, Google Trends is also available at higher frequencies, but the data start only in 2004.

A drawback of our news-based indices is that they are potentially subject to the same measurement problems as clinical measures as the latter are an important input into the media reports that we exploit. For example, there are different testing procedures and manifold reporting practices both over time and across countries. Furthermore, typically there is a time lag between the outbreak of an epidemic and the reported cases and associated deaths. In addition, the indices can capture reporting about the economic effects of a disease. Finally, the media coverage might contain some noise reflecting hypes or underreporting. Therefore, we next conduct a careful validation of the news-based indices with those clinical measures that are available for an overlaping sample. Moreover, in the sensitivity analysis, we show that the results hold when accounting for measurement error and when cleansing the index of reporting about the economic damages of epidemics.

2.3. Validation of news-based viral disease indices

To provide a first visual validation of the monthly indices, the shaded areas in Figure 1 show the retrospective dating of the epidemic or pandemic by the WHO. The areas cover the main spikes of the indices, suggesting that the latter track the start and end points of the official definitions and, hence, the timing of epidemics closely. To see whether the height of the indices also provide a reasonable approximation of the severity of the diseases, we perform a formal external validation. We cannot directly test the SARS, MERS, and swine flu indices as monthly data on global infected or fatalities are not available for these diseases. But we can externally validate the two influenza indices using national data on the monthly number of total (excess) deaths in Germany and the U.S.

Table 1 shows regression results. The dependent variable in columns 1-3 is the influenza index for Germany, while in columns 4-6 it is the influenza index for the U.S. Thereby, we test whether the media capture the true extent of the disease well. We could also reverse the regressions as we are here only interested in partial correlations and not causality. For both indices, the first two columns estimate models in levels, while the last column refers to a model in differences. In all models, we use robust standard errors and standardize the dependent and the main explanatory variable to simplify the economic interpretation within and across models. Moreover, all models contain the lagged endogenous variable, year dummies and month-of-the-year effects to control for potential autocorrelation, secular trends, and seasonal patterns, respectively. For Germany, the sample starts in 1990M1 and for the U.S. in 1998M1 due to data availability. Generally, the models describe the data well. The R^2s are between 17% and 67%.

Dependent variable	Influenza index Genios			Influenza index NYT		
	level	level	change	level	level	change
	(1)	(2)	(3)	(4)	(5)	(6)
Lagged dependent variable	0.43***	0.43***	-0.19	0.40***	0.41***	-0.11
	(0.12)	(0.11)	(0.16)	(0.14)	(0.14)	(0.13)
Total deaths Germany	0.31**					
	(0.13)					
Excess deaths Germany		0.16**				
		(0.07)				
Change total deaths Germany			0.47***			
			(0.08)			
Total deaths U.S.				0.59***		
				(0.19)		
Excess deaths U.S.					0.20***	
					(0.07)	
Change total deaths U.S.						0.75***
						(0.20)
Break point dummy	-0.12	-0.10	-0.14	0.06	-0.05	0.05
	(0.09)	(0.10)	(0.08)	(0.10)	(0.11)	(0.12)
Month effects	yes	yes	yes	yes	yes	yes
Year effects	yes	yes	yes	yes	yes	yes
Monthly observations	361	361	360	255	255	254
R^2	0.67	0.66	0.29	0.59	0.59	0.17

Table 1: Regression analysis of monthly news-based influenza indices. *Notes:* The table shows results for alternative models that regress the news-based influenza indicators based on text analysis of the Genios data (column 1-3) and the NYT media archive (column 3-6) on external measures of disease intensity and control variables. Robust standard errors are in parentheses. *,**,*** denote statistical significance at the 10%, 5% and 1% level, respectively.

In column 1, the main explanatory variable is the number of deaths in Germany. It is strongly positively correlated with the influenza index. The relation is significant at the 5% level. In column 2, we use the number of excess deaths in Germany to allow for the possibility that the month and year dummies do not fully capture seasonal patterns and secular dynamics due to changing demographics. To compute the excess deaths, we fit a higher-order polynomial for each month to the number of total deaths and subtract it from the latter. The point estimate drops in size as we remove some of the lower frequency correlation between the index and the measure of mortality, but the coefficient remains significant. In column 3, we return to total deaths but estimate a model in differences to treat potential secular trends in the index and mortality symmetrically. The point estimate increases to 0.47 and is now highly significant with a standard error of 0.08.

The results are similar for the U.S. in columns 4-6. The point estimates for the main explanatory variable are larger than in columns 1-3 and more precisely estimated. They are all statistically significant at the 1% level. They suggest that a one standard deviation increase in the explanatory variable leads to an increase in the influenza index by between 0.2 and 0.8 standard deviations. Column 6 suggests a nearly one-to-one mapping between total deaths and the influenza index. The null hypothesis that the point estimate is one is not rejected. The p-value of the Wald test is 0.21.

To see whether the relations are stable over time, the models allow for a potential break. In columns 1-6, we interact each mortality measure with an indicator variable for the period post 2004. In this way, we allow for different coefficients before and after SARS. There is no clear sign pattern for the interaction terms and none of them is significant. The relation between the mortality measures and the media indices seem to be stable. In other words, there is no evidence of time-varying media coverage of diseases.

Finally, we assess whether the media give appropriate weights to the clinical characteristics of an epidemic. We can do this for the current SARS-CoV-2 outbreak for which monthly data on global Covid-19 cases and deaths are available since 2020M1. As dependent variable, we employ the Covid-19 or coronavirus news index. The indices are correlated with 0.84. As clinical measures, we use the log number of global Covid-19 deaths and cases. As before, we standardize both the dependent and explanatory variable. We estimate all models in first differences to account for common trends. Furthermore, to save degree of freedoms in the short sample, we do not include the lagged endogenous variable to control for autocorrelation but use HAC standard errors with four lags. All models include a constant.

Table 2 shows the regression results. In columns 1 and 2, the dependent variable is

Dependent variable:	Covid-19 index		Coronavirus index	
	(1)	(2)	(3)	(4)
Global deaths Covid-19	1.97**		2.14**	
	(0.75)		(0.97)	
Global cases Covid-19		2.13**		2.28**
		(0.84)		(1.05)
Observations	21	21	21	21
F-statistic regression	6.94	6.42	4.88	4.72

Table 2: Regression analysis of news-based coronavirus indices. *Notes:* The table shows results for regressions of the Covid-19 index (columns 1-2) and the coronavirus index (columns 3-4) on global Covid-19 deaths and cases, respectively. The indicators are based on text analysis of the Genios media archive. All models include a constant. All variables are standardized and in first differences. HAC robust standard errors for 4 lags are in parentheses. *,**,*** denote statistical significance at the 10%, 5% and 1% level, respectively.

the Covid-19 index. In columns 3 and 4, it is the coronavirus index. We have only 21 observations but all point estimates are significantly positive at the 5% level. The models' fit is decent with F-statistics of 5-7. The estimated elasticities indicate some over-reporting of the media. The coefficients are around 2. This number suggests that a one standard deviation increase in cases is associated with a two standard deviations increase in articles about the disease. Overall, we conclude that the disease indices are valid approximations of the timing and intensity of regular influenza, epidemics, and pandemics.

3. The dynamic impact of epidemic shocks

The derivation and presentation of the core results proceeds in three steps. In Section 3.1, we present the conceptual and empirical model. In Section 3.2, we analyze major epidemic shocks in the base sample individually and contrast them with seasonal influenza shocks. In Section 3.3, we construct an aggregate news-based disease index, estimate the average effects of epidemic shocks on the world economy, and compare these estimates to the impact of Covid-19 out-of-sample.

3.1. Conceptual and empirical framework

Viral disease outbreaks can affect economic activity through a number of transmission channels. On the supply-side, there is direct absenteeism from work due to mortality, people who die, and morbidity, those who are infected and/ or are in quarantine. Quarantines and lockdowns are likely to reduce output despite new technologies facilitating online collaboration. There might also be prophylactic and indirect absenteeism because people avoid going to work where they might become infected, because they need to care for others who are ill, or because schools are shuttered. Public authorities might also close workplaces. A reduction of business trips can lead to fewer contracts and, thus, fewer orders. The closing of borders disrupts international value chains and falls in production in certain regions lead to negative supply chain shocks worldwide, which result in decreased international trade and output.

On the demand-side, there are direct and indirect effects as well. Due to mortality and morbidity, consumption demand will decline immediately. The closing of workplaces, shops, and more general curfews lower actual consumption of products of the affected branches and sectors. Losses of social consumption (tourism, going to restaurants, attending public or social events) are likely to be permanent. Moreover, consumers and investors wait-and-see when faced with higher uncertainty, purchases of durable goods and other investment decisions could be postponed. Financial investors are likely to reduce exposure to risky assets. Finally, a crucial determinant of the overall impact of epidemics is the response of fiscal policy, which might aim at stabilizing demand.

To identify and trace out the economic impact of epidemic shocks, we use the following vector autoregression (VAR) for a monthly frequency:

$$y_t = c + \Pi(L)y_{t-1} + u_t.$$

The $k \times 1$ vector c includes constant terms, the matrix $\Pi(L)$ in lag polynomials captures the autoregressive part of the model, and the vector u_t contains k serially uncorrelated reduced form shocks with $u_t \sim N(0, \Sigma)$.

In the baseline specification, the set of variables in y_t includes a news-based disease index and the logarithm of industrial production in China, South Korea, Germany, France, Canada, the U.S., and the world, respectively. We focus on those countries to cover parts of Asia, Europe, and North America explicitly, while keeping the model estimable. Moreover, China and South Korea had many SARS cases, and swine flu was widespread in Canada and the U.S. We use industrial production because it captures economic activity, is a standard measure of the business cycle (Stock and Watson, 1999), and is available at the monthly frequency for this set of countries (unlike services or employment).

In alternative specifications, we keep a fixed set of baseline variables in y_t and rotate other variables of interest in, one at a time. For example, we rotate in measures of inflation, employment, or fiscal variables. This approach follows Ramey (2011) and is a particularly flexible. It does not require a Bayesian perspective, a panel VAR, or factor structure to deal with the curse of dimensionality as the baseline model already contains $k(kp+1)+\frac{k(k-1)}{2}$ paramters, where p is the lag length. We show that the results are robust to using local projections (Figure A.9).

The usual lag length selection criteria suggest 1-3 lags. However, as these are known to underestimate the true lag length, we set $p = 6$ to obtain reliable predictions for the annual horizon. The results are similar when changing the lag length (Figure A.7). The typical sample is 1991M1-2019M12 when using disease indicators based on the Genios archive and analyzing the world economy. The sample changes somewhat across specifications, depending on the variables included. We stop in 2019M12 to exclude the corona pandemic. Including it would dwarf all other outbreaks (Figure 1) and essentially imply a dummy variable regression of production on the early months 2020, simply capturing the observable output losses in these months. Instead, we perform two out-of-sample analyses. First, at the end of this section, we assess how well the model predicts the effects of the Covid-19 pandemic in 2020. Second, in the next section, we use the news index for the U.S. It is available since 1919M1, thus including further pandemics. Online Appendix A lists the definition, construction, and sources of the variables.

The innovations u_t are assumed to be linearly driven by an epidemic shock ϵ_t^e, which we aim to identify, and other structural shocks ϵ_t^*, which are of no interest for this paper:

$$u_t = b^e \epsilon_t^e + B^* \epsilon_t^*.$$

The $k \times 1$ vector b^e captures the impulse vector to an epidemic shock of size 1. To identify this shock, we rely on a Choleski decomposition of $\Sigma = BB'$, with $B = [b^e, B^*]$ a lower triangular matrix and where we have normalized the variances of the structural shocks to one, $\epsilon_t \sim N(0, I_k)$. Inference is based on a standard fixed-design residual wild bootstrap with 1000 replications.

We order the disease index first to identify economy-wide health shocks. The key identifying assumption is that the news-based index is contemporaneously exogenous to economic activity at the monthly frequency. This ordering is standard when news events are directly observed or measured through textual analysis and included in structural VARs (Ramey, 2011; Mertens and Ravn, 2013). In our case, it implies that news about viral disease outbreaks are determined independent of current production levels. As we are only interested in the health shock, the ordering of the remaining variables is irrelevant.

In public economics, there is no consensus about the effect of economic activity on health status. While in the long-run there seems to be a positive relationship between GDP per capita and public health (Pritchett and Summers, 1996), others argue that mortality decreases during recessions because smoking and obesity declines (Ruhm, 2000) or because fewer people are traveling, thus reducing interpersonal contact and the spreading of infectious diseases (Adda, 2016). If the latter two arguments hold at the monthly frequency, our estimates would reflect a conservative estimate of the adverse impact of viral disease outbreaks on economic activity. Nevertheless, in Section 5, we show that endogeneity is unlikely to affect our results. We use a higher frequency version of the disease index, a cleaned version that controls for media reporting about the economic effects of epidemics, and an alternative ordering of the variables (Figure 9).

Other potential source of endogeneity are omitted variables and measurement error in the disease index. Section 5 shows that theses are also unlikely to drive our results. We estimate factor-augmented models following Bernanke et al. (2005) (Figure 10) or add a large number of economic variables to the model (Figures A.2, A.3). We also construct and include an index of containment policy to control for health interventions (Figure A.5). Furthermore, we use the instrumental variable approach of Stock and Watson (2012) and Mertens and Ravn (2013) that accounts for various forms of errors-in-variables (Figure 11). Finally, pulmonary infections can have seasonality, in particular regular influenza. Therefore, we document the robustness of the estimates to including month-of-the-year fixed effects (Figure A.6, A.8).

3.2. The impact of SARS, swine flu, and seasonal influenza

First, we analyze two major adverse health shocks in the base sample separately. According to the maximum index values, these are the SARS epidemic of 2002/03 and the swine flu pandemic of 2009/10 (Figure 1). The estimated dynamic effects on production are shown in the first two columns of Figure 2. Each column is based on a separate model where we order either the SARS index or the swine flu index first, while keeping the other variables constant for comparison. Each shock is scaled such that the maximum index response corresponds to the peak value of the index in the sample. Thus, the responses measure the estimated economic damage of each viral disease outbreak.

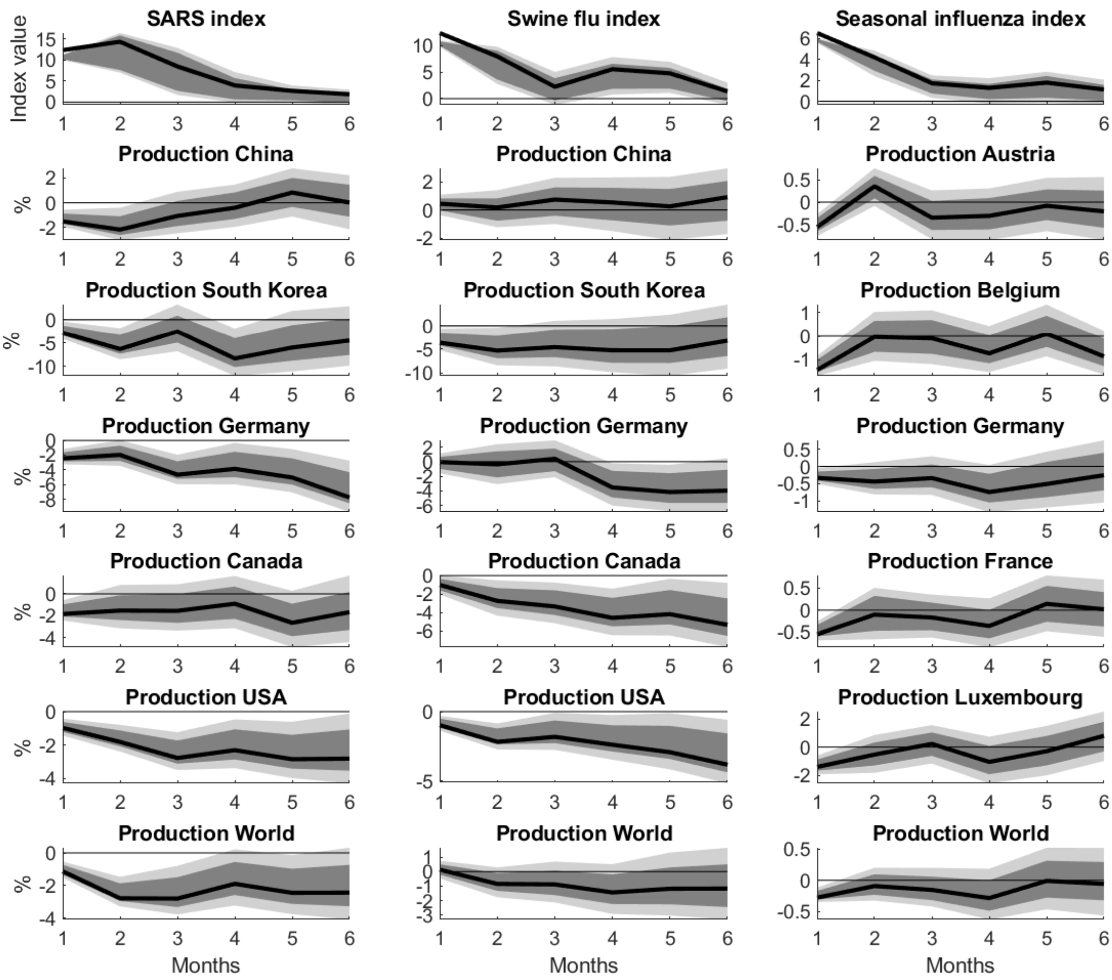

Figure 2: SARS, swine flu, and seasonal influenza shocks. *Notes:* The figure shows the dynamic impact of a SARS shock (column 1), a swine flu shock (column 2), and a seasonal influenza shock (column 3), obtained from three structural VAR models, on production in single countries and globally over 6 months, along with 68% and 90% bootstrapped confidence bands.

Both shocks entail significant adverse effects on output in major economies and around the world as a whole. The SARS shock led to a significant drop of industrial production in China by a maximum of about 2%. Output recovers within two quarters. As the quality of Chinese data is unclear and because production data for Hong Kong, another epicenter of the epidemic, are not available, we also look at production in South Korea, which is a main trading partner of China. The impact of the SARS shock indeed appears stronger according to this metric. Economic activity falls by up to 10% one quarter after the shock. For Germany, the U.S., and Canada, the effects are successively smaller and more delayed. World output falls significantly by more than 2% for two quarters.

The estimated effect of the swine flu outbreak is by and large similar, with several exceptions. Output in China is not significantly affected, while production in Canada and the U.S. drops by more than in the case of SARS. This is plausible as both countries where, along with Mexico, the epicenter of this pandemic. The impact on production in South Korea is roughly similar to the effect of the SARS shock. For Germany, there is a delayed response. Global output declines by more that 1%, but the impact is only borderline significant at the 68% confidence level.

The last column estimates the impact of seasonal influenza shocks on production. We use the flu index for Germany (Figure 1), but winsorize the series at the 95th percentile to chop the spikes related to avian and swine flu and to concentrate on the typical influenza effects.[4] There are nine non-epidemic peaks in the range between 4 and 6.5 in the series. We scale the shock to the largest of these values, that is, to 6.5, which corresponds to January 2000. As the disease index is constructed from German-speaking media, we include output of Austria, output of a set of neighboring countries of Austria and Germany, which have roughly similar climate and hence influenza courses, and world production into the model. The disease index is significantly positive for about one quarter, reflecting the typical duration of the influenza season. Production falls significantly in all countries individually and globally. In the single countries, the fall is between 0.5% and 1% and one-off. Production returns to its initial level already in the second month after the shock.

For the world as a whole, the drop is 0.2%. This magnitude seems plausible as Europe accounts for a significant fraction of world production and since influenza cycles are correlated across the Northern hemisphere. Comparing the impact of the influenza shock to the average effect of the SARS and the swine flu shock shows that the latter two are much more damaging. Their cumulative output loss, that is, the area between the zero line and the point estimate, is on average 19 times larger for the first half year alone. Still, the significant impact of seasonal influenza may raise the question why vaccination against it is not compulsory. One reason might be that this is costly, too. Another could be that there is a strong public opposition against compulsory vaccination in many countries.

[4]An alternative way to focus on the impact of regular influenza is to use the index as it is and to control for the flu pandemics by including corresponding indices as exogenous variables and letting the latter pick up the impact of the pandemics. The results are essentially the same.

3.3. The global economic impact of epidemic shocks

Now, we study the average effects and international propagation of epidemic shocks. For this, we construct an aggregate epidemic index by summing over the indices for SARS, MERS, and swine flu for the period 1991M1-2019M12. We focus on these subindices as the underlying viral diseases are comparable in terms of transmissability and mortality (Section 2.1). As mentioned, we end the sample in 2019M12 and disregard the Covid-19 index as the corresponding observations would dominate all historical episodes.

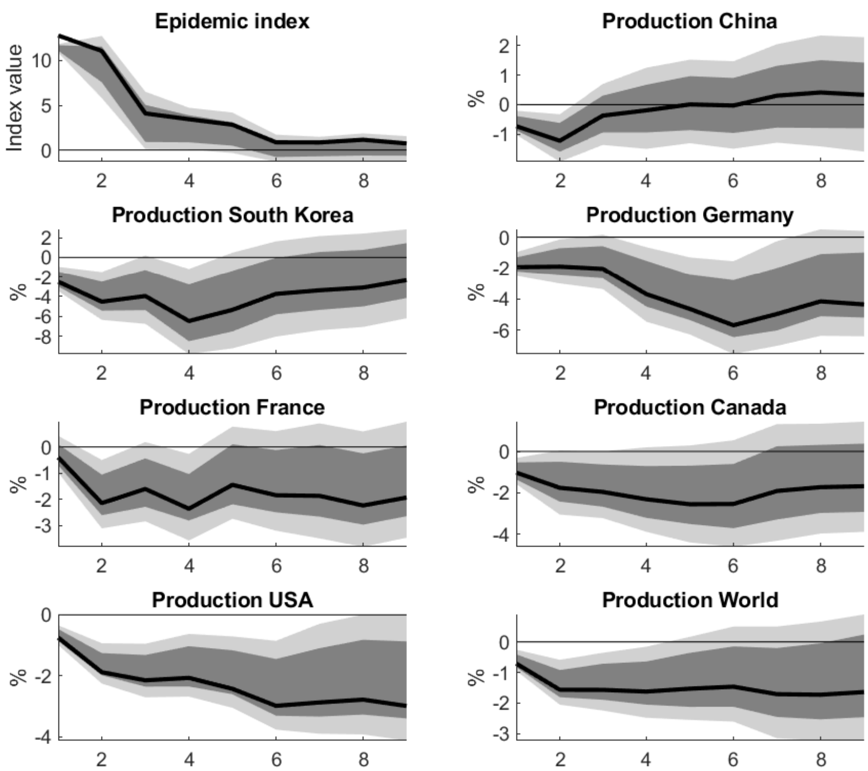

Figure 3: Global output effects of epidemic shocks. *Notes:* The figure shows the response of industrial production to an epidemic shock of size 12.8 over 9 months along with 68% and 90% confidence bands.

We scale the shock to the maximum of the aggregate epidemic index in the sample, which is 12.8 and corresponds to the height of the SARS epidemic in 2003M4. The other peak corresponds to the swine flu pandemic and is comparable, with value of 12.5 in 2009M11. Figure 3 shows the results. They are similar to (an average over) the first two columns of Figure 2. Output falls in all considered countries and the decline is mostly significantly different from zero. Except for China, the drop lasts for at least two quarters according to the point estimates, before output convergence back to trend. World production falls by close to 2% in the first three quarters following the shock. The effect is persistent. There is no recovery at the end of the horizon. According to the 68% bands, overshooting can be ruled out. The output losses seem permanent. This finding adds to the evidence of Cerra and Saxena (2008), who show that the output costs of financial and political crises, ranging between 4-12% over several years, are typically not recuperated.

To analyze the international propagation of epidemic shocks, we use the baseline model with one modification. We replace industrial production in France by one alternative global variable of interest at a time. Replacing output in France is somewhat arbitrary, but replacing production of any of the other smaller economies yields similar conclusions. Figure 4 shows the results. The arrival of foreigners at airports drops upon impact by 5%. The decline is statistically significant for more than a quarter, but even after three quarters the fall is not fully redeemed. In contrast, airport arrivals of residents increase sharply one month after the shock as people fly home. Imports of emerging market economies drop by about 3%, possibly reflecting shuttered factories requiring fewer inputs, but the effect is short-lived. It is passed-through to exports of emerging market economies with a delay of about one quarter. Exports of advanced economies drop immediately and significantly, while their imports respond to the shock with some lag but then decline by roughly the same amount and similarly persistently. Overall, world trade falls significantly by about 2% for more than two quarters.

Following declines in production and trade, world employment falls by 0.5% two quarters after the shock. It recovers only gradually. In contrast, world retail sales increase significantly for several months, probably reflecting panic buying. However, world equity prices, measured by the MSCI world index, fall upon impact and decline further, reflecting the overall negative economic consequences of epidemics. Consistently, global consumer prices decline significantly. They show similar dynamics as employment with at trough response of 0.4% after six months and slowly returning towards the level where they would have been without the shock after three quarters. The negative demand effects depress consumer prices rather in advanced than in developing economies, as the next three panels show. In line with the global price drop, the Federal Reserve lowers the policy rate by about 50 basis points.

Finally, we confront the conditional predictions of the model with actual data on the Covid-19 pandemic in an out-of-sample analysis to assess the model's performance and to document commonalities with and differences to the historical viral disease outbreaks. We

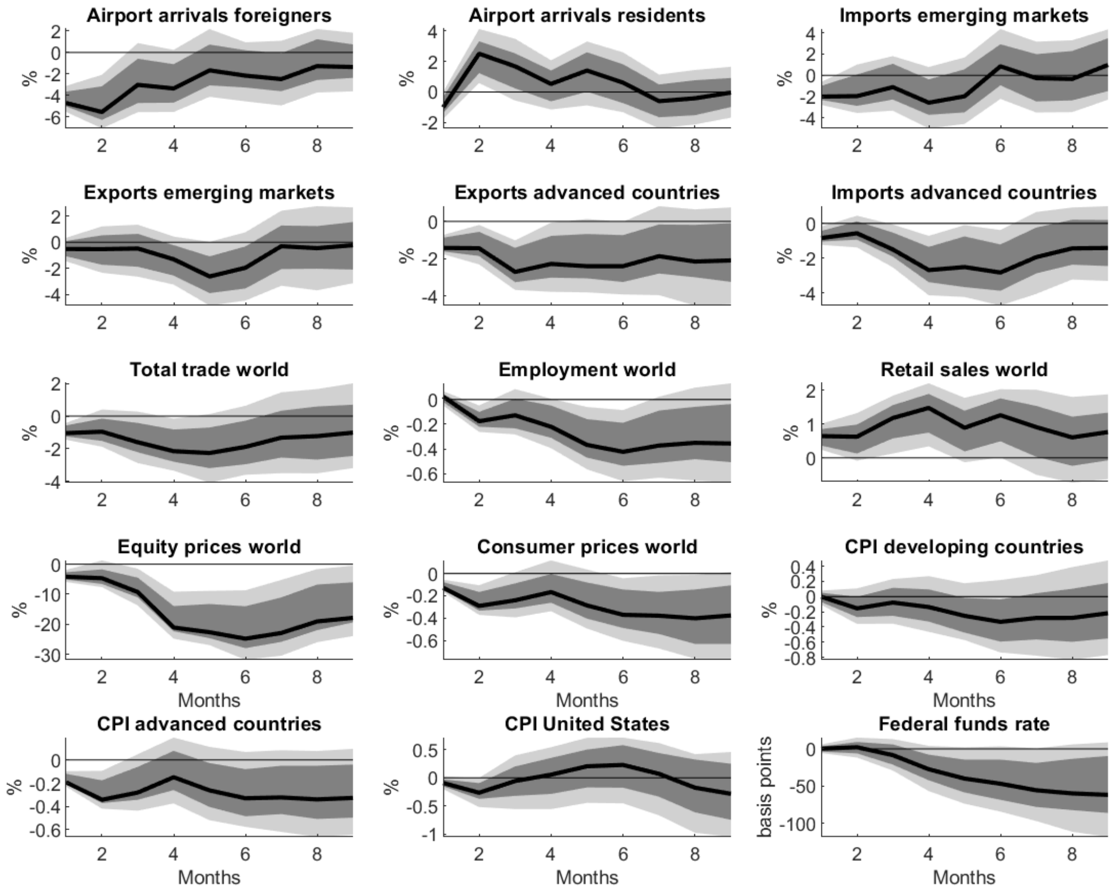

Figure 4: The international propagation of epidemic shocks. *Notes:* The figure shows the response of global variables to an epidemic shock of 12.8 over 9 months along with 68% and 90% confidence bands.

use the estimated model of Figure 4 and feed it with two shocks. The thought experiment is that the start of the Covid-19 pandemic in 2020M2-M3 was unexpected. Hence, we scale the two shocks to the change of the corresponding news indices in these two months (Figure 1). As we have two indices, we take the mean of the coronavirus and Covid-19 index to avoid double counting. The changes in the mean index are 18.0 and 104.8 in these two months. Of course, there is large uncertainty about the calibration of the shocks. We compute responses for 9 and 8 months for the first and second shock, respectively. Then, we add the responses (and error bands) for the second shock to those for the first shock from period 2 onward. Finally, we compare the summed responses to the actual data on the mean index and industrial production for the period 2020M2-M11. For production, we compute the percentage output loss relative to mean output in 2019Q4 as log differences.

Figure 5 shows the results. Compared to the previous figure, the effects are much larger (in absolute value) as the shocks are larger. The responses are also more persistent

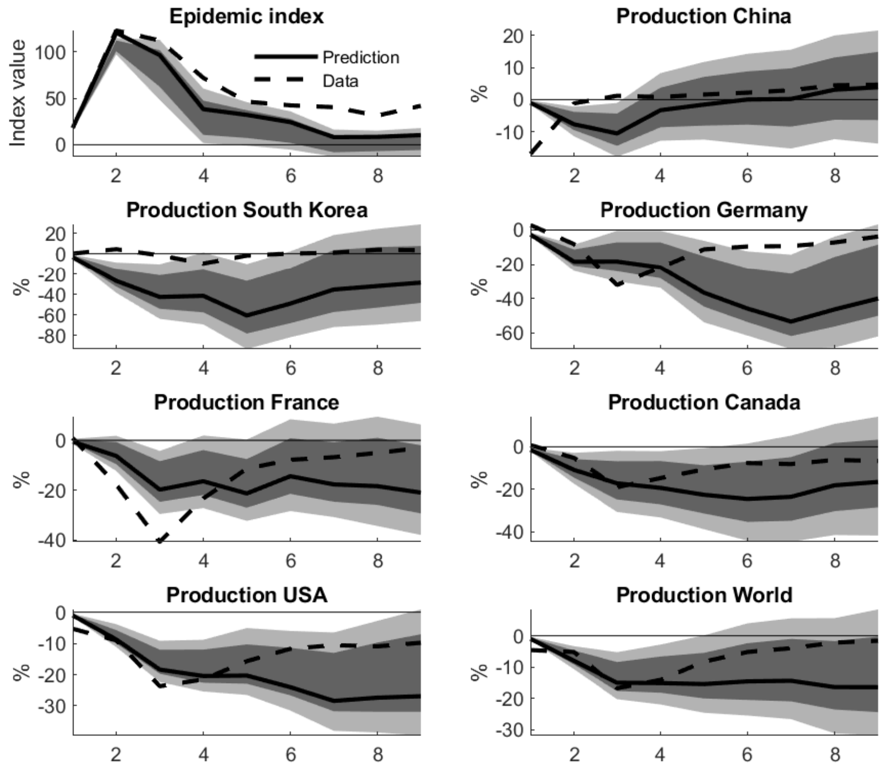

Figure 5: Out-of-sample prediction for Covid-19. *Notes:* The solid lines show the summed responses of the epidemic index and production to one epidemic shock of size 18.0 at $h = 1$ and one epidemic shock of size 104.8 at $h = 2$, along with 68% and 90% confidence bands. The dashed lines refer to actual data. Production losses are computed relative to 2019Q4.

because now there are two shocks, one at $h = 1$ and one at $h = 2$. The simulated dynamics of the epidemic index match the data closely for the first two months by construction. Then, the model somewhat underpredicts the persistence of the Covid-19 pandemic.

The simulated paths for the production losses cover by and large the actual trajectories. The patterns for China and South Korea differ from those for the other countries, with the model having more difficulty in matching the data. For China, the actual production drop is earlier since the Covid-19 pandemic has hit China already in 2020M1. For South Korea, the model clearly overpredicts the losses. One interpretation is that the country has learned from the SARS outbreak how to successfully mitigate Covid-19. For the other countries and for the world as a whole, the model does a surprisingly good job in matching the size of the output losses, although the model tends to overpredict the effects. Interestingly, the data trough earlier and return to trend quicker. Potential reasons are the drastic containment measures and large fiscal stimuli that were implemented around the globe to first mitigate the pandemic and then help the economic recovery from it.

4. Country studies

In this section, we study the impact of epidemics on individual countries to obtain an impression of the direct and indirect effects of epidemic shocks (Section 4.1). Then, we conduct a historical analysis based on the longer sample for the U.S. to validate and extend the results for the world economy based on the more recent sample (Section 4.2).

4.1. Direct and indirect effects

First, we look at Hong Kong, which was most affected by SARS in terms of the clinical attack rate and mortality rate. Moreover, it provides more detailed economic data at the monthly frequency than China or Taiwan, the other two epicenters of SARS, allowing for a more granular view. We take these estimates as an approximation of the direct economic effects of epidemics. We contrast these findings with estimates for Germany. This is a small open economy heavily relying on international value chains. At the same time, it had low attack and mortality rates for both SARS and swine flu, such that the largest effects on the German economy arguably occurred through indirect trade and confidence effects. Moreover, the country provides rich economic data at the monthly frequency.

Figure 6 shows the results for Hong Kong. They are based on the model in the first column of Figure 2 with the SARS index, and adding one additional Hong Kong-specific variable at a time. As before, the shock is scaled to the maximum of the SARS index. In response to the epidemic shock, airport arrivals essentially come to a halt. It takes six months before they return to their pre-shock trend. Airport departures drop by 60%, also needing two quarters before recovering. Business and tourism travel also take a hard hit. The hotel room occupancy rate drops by 60 percentage points upon impact and by nearly 80 percentage points after two months. In contrast, exports and imports are only mildly affected as the manufacturing sector is outside the city center in Pearl River Delta and it largely relies on immigrant workers with few potentially infectious ties to the main community. Unfortunately, there are no production data for Hong Kong at the monthly frequency. Nevertheless, the significant fall of employment by more than 2% suggests that the overall effect of the shock on the economy is severely negative.

Indeed, retail sales fall by 10% and for more than one quarter. The strong decline of demand for travel, tourism, and retail goods is reflected in a sustained fall of consumer prices by more than 2%. The collapse of economic activity is also associated with a

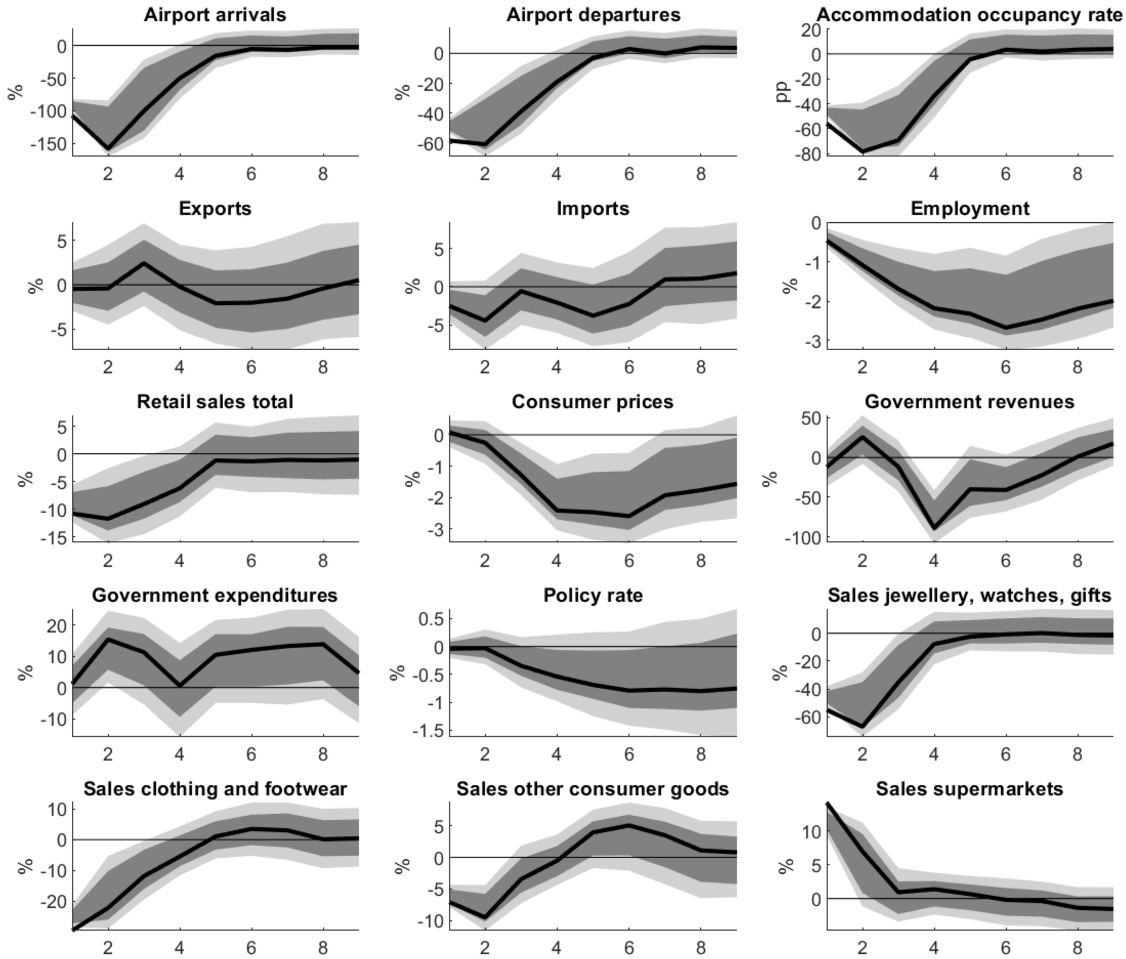

Figure 6: Impact of SARS shock on Hong Kong. *Notes:* The figure shows the response to a SARS shock of size 11 over 9 months of Hong Kong-specific variables individually added to the SVAR of the first column of Figure 2, along with 68% and 90% confidence bands.

pronounced, lagged fall of government revenues. Public expenditure policy and monetary policy are expansionary but cannot offset the demand contraction. Finally, when looking at selected components of domestic trade to see which sectors suffer most, we observe the largest declines for retail sales of jewellery, watches, and valuable gifts (–60%), clothing, footwear, and allied products (–30%) and sales of other consumer goods (–10%). The typical domestic trade contraction lasts for about one quarter and the foregone sales are largely permanently lost. Only for other consumer goods is there some indication of catching-up. The notable exception from these patterns are supermarket sales. They increase drastically by 14% as consumers stockpile necessities, substitute restaurant visits, and more generally refrain from social consumption activities outside their homes.

Next, we analyze the impact of epidemic shocks on Germany to obtain a quantitative impression of their indirect effects. We return to the aggregate index and add to the model one Germany-specific variable at a time. The responses of the latter are shown in Figure 6. The shock is scaled to the maximum of the epidemic index (12.8). The composition of variables differs somewhat from the selection for Hong Kong as more economic data series for the monthly frequency are available. But we also include employment to compare the

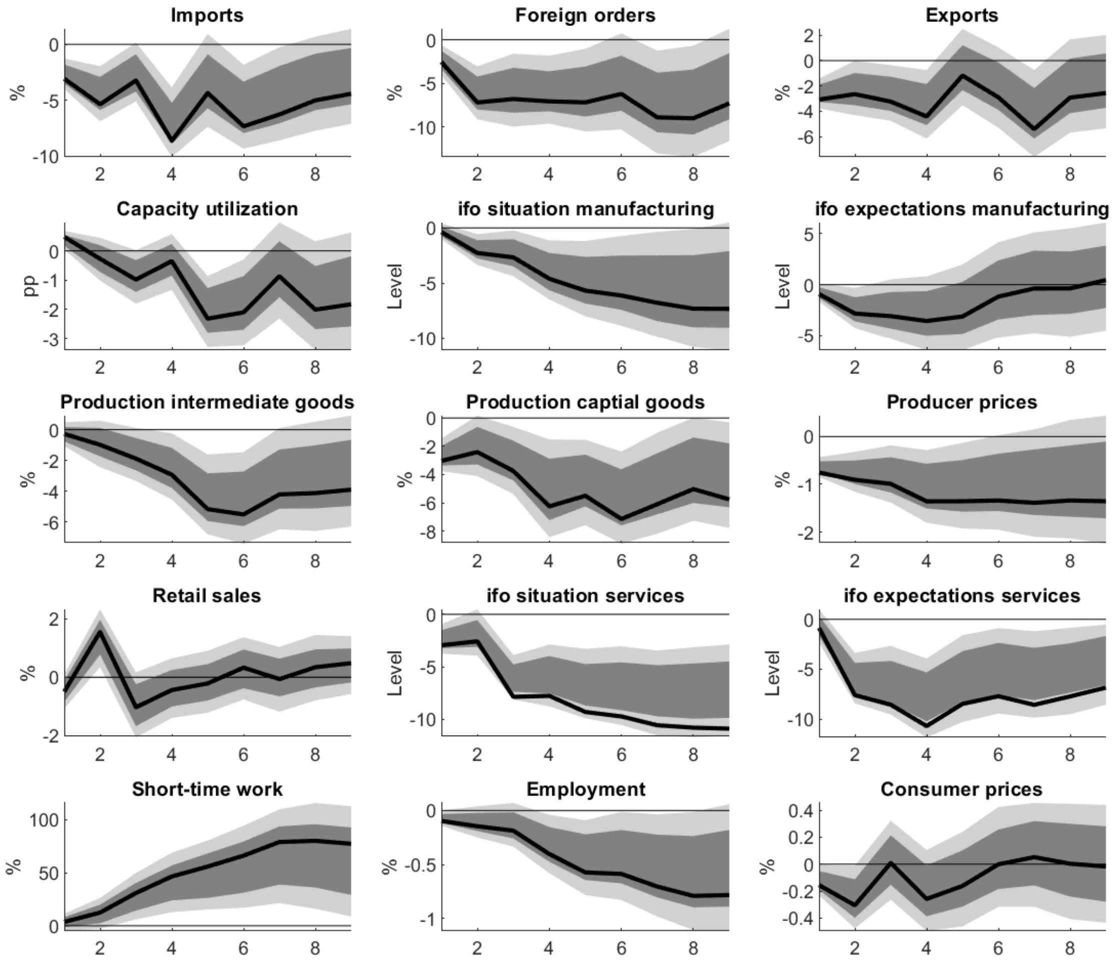

Figure 7: Impact of epidemic shock on Germany. *Notes:* The figure shows the response to an epidemic shock of size 12.8 over 9 months of Germany-specific variables individually replacing production in France, one at a time, in the baseline SVAR underlying Figure 3, along with 68% and 90% confidence bands.

The figure shows that imports decline significantly upon impact and subsequently fall further as inputs into German supply chains are not shipped. Foreign orders also drop immediately and by similar amounts. Exports, on the other hand, fall less, perhaps because companies can partially replace imported inputs by domestic products, either out of their stocks or from home companies. However, substitution seems imperfect and capacity utilization declines. More generally, firms are more pessimistic about their current and expected future situation, as judged by the ifo business survey for the manufacturing sector. Looking at the subcategories of production shows that, in particular, intermediate and capital goods suffer, potentially due to a disruption of international supply chains. Overall, producer prices fall by 1% in response to the lower demand for their goods.

As for Hong Kong, the decline in economic activity is not universal. Retail sales increase significantly one month after the shock, potentially reflecting panic buying and consumption substitution. However, in general, the service sector is also persistently negatively affected by the shock, as the corresponding decline in the ifo index current situation and in the expectation component show. Hence, overall, short-time work increases strongly (although from a typically very low level) and employment falls significantly by more than 0.5%. Comparing the employment decline to that of Hong Kong suggests that the indirect impact of epidemic shocks is about one-fourth of the direct effect. Finally, the economic contraction in Germany is associated with lower consumer prices.

4.2. Historical analysis of epidemic shocks for the U.S.

The last part of the main analysis is an estimation of the average impact of epidemic shocks on the U.S. economy using historical data. This serves two purposes. First, it is an external validation of the previous results for the world economy, which are based on news-indices from the Genios media data, using an alternative news data source, the NYT archive. Second, the NYT data extend further back in time and, hence, contribute several (global) influenza outbreaks that are interesting in their own right to the sample. Specifically, the data cover the epidemic influenza of 1929, the 1957/58 pandemic Asian flu (H2N2 virus), and the 1968/69 pandemic Hong Kong flu (H3N2 virus) (Figure 1).

To exploit as many of these additional observations as possible, we only include variables in the model that are available at the monthly frequency at least from before the first of these episodes, that is, from before the influenza pandemic of 1929. Hence, the model includes the influenza index, the logarithm of, respectively, industrial production, oil and gas production, the S&P 500, and consumer prices; as well as the three-month rate on AA-rated commercial papers as a proxy for monetary policy. As before, the model contains six lags and we scale the shock to the maximum of the flu index for the U.S., which is 41 and corresponds to the peak of the pandemic influenza in May 2009. The peaks of the other flu outbreaks are between 27 and 35, and, as such, roughly comparable.

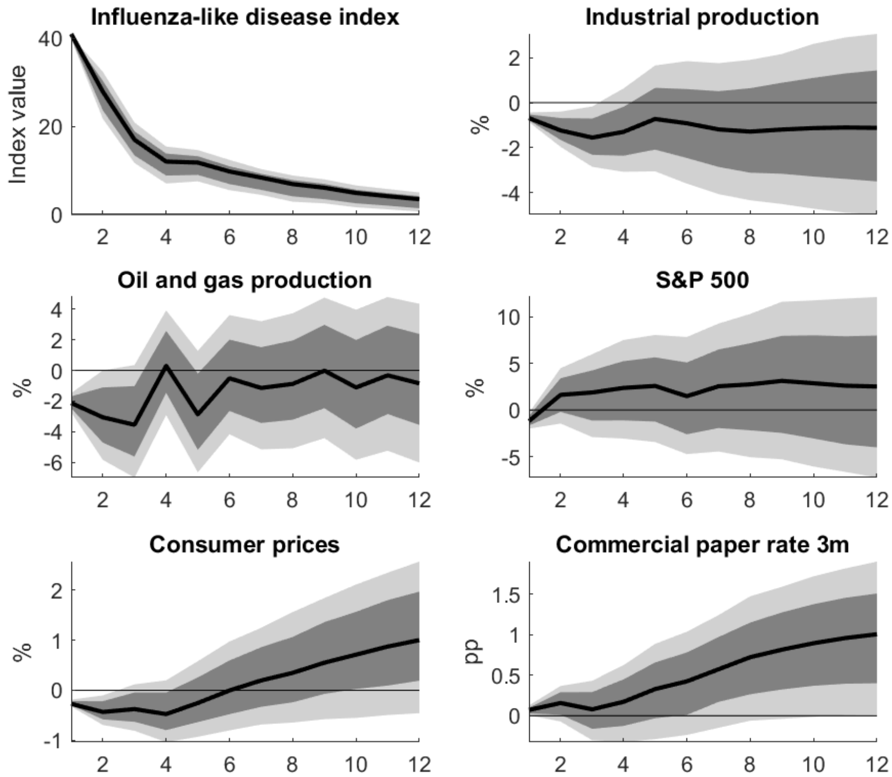

Figure 8: The impact of influenza-like disease shocks on the U.S. economy. *Notes:* The figure shows the responses of the endogenous variables to an influenza-like disease shock of size 41 over 12 months based on a SVAR for the U.S., along with 68% and 90% confidence bands.

Figure 8 shows that the influenza index increases significantly upon impact and remains elevated for one year. This suggests that the identified shocks largely reflect pandemic or epidemic flu outbreaks, that is, the outliers in Figure 1, rather than unexpectedly severe regular influenza. While the former typically last 6 to 18 months, the latter are short-lived and confined to the winter months. In response to the disease shock, industrial production declines significantly upon impact and falls further for two months. Thereafter, it returns to trend. The trough is –1.6%. Oil and gas production also falls upon impact, by more than industrial production. It also bottoms one quarter after the shock and then recovers. Equity prices drop significantly by 1.2% and then overshoot slightly. Consumer prices decline by 0.2% initially. They fall further, to –0.4%. After two quarters, there is evidence of overshooting as well. Finally, somewhat surprisingly given the disinflation, the three-month rate tends to increase (although the response is largely insignificant at the 90% level). This positive reaction can mirror an anticipation of monetary policy of the future price increase.

5. Sensitivity analysis

We perform an extensive sensitivity analysis. Here, we focus on endogeneity concerns. We perform more tests that we summarize at the end of the section. For exploring whether the estimates are affected by reverse causality, we now conduct three analyses. First, we modify the epidemic index such that it refers to the first week of each month only to minimize the overlap with the monthly production data. As the news-data are available at the daily frequency, we construct the subindices entering the aggregate index using keyword counts in each onset week of the month. Then, we sum these first-week subindices within months to obtain an aggregate index that disregards media coverage of epidemics of weeks 2-4. This implies that the contemporaneous overlap between the modified epidemic index and industrial production is only 25%. Thus, the identification strategy assumes that the news-based disease index is not caused by weekly production. The dashed lines in Figure 9 show that the main results hold. The point estimates are close to the baseline ones (solid lines) and lie within the 68% confidence bands for those.

As second test, we cleans the epidemic index to address concerns that it might pick up reporting on the economic effects of epidemics. If this was the case, we would underestimate the impact of epidemic shocks as the output losses in the data during epidemics would be associated to lower index values. For each subindex we count the joint occurrence of the respective disease word and 'economic/economy' (Wirtschaft in German).[5] We scale the subindices as before and subtract them from the scaled aggregate index. Figure A.1 shows the cleaned index spikes at the same time as the baseline index, while the peaks are somewhat lower. The latter two observation indicates that indeed some news on epidemics is associated to their economic consequences. The dotted lines in Figure 9 show that the estimated impact of epidemic shocks increases when using the cleansed index. The point estimates are largely within the confidence bands for the baseline.

As third check, we return to the baseline epidemic index but reorder the variables in the model. We place world production before the index and identify the model recursively, as before. Now, we focus on the second shock, which is orthogonal to current world

[5]Performing the same query for the joint occurrence of the disease word with either 'unemployment' (Arbeitslosigkeit), 'income' (Einkommen), or 'recession' (Rezession) yields substantially fewer hits such that we focus on the joint occurrence with 'economic/economy', which provides an upper bound for the impact of cleaning the index on the estimates.

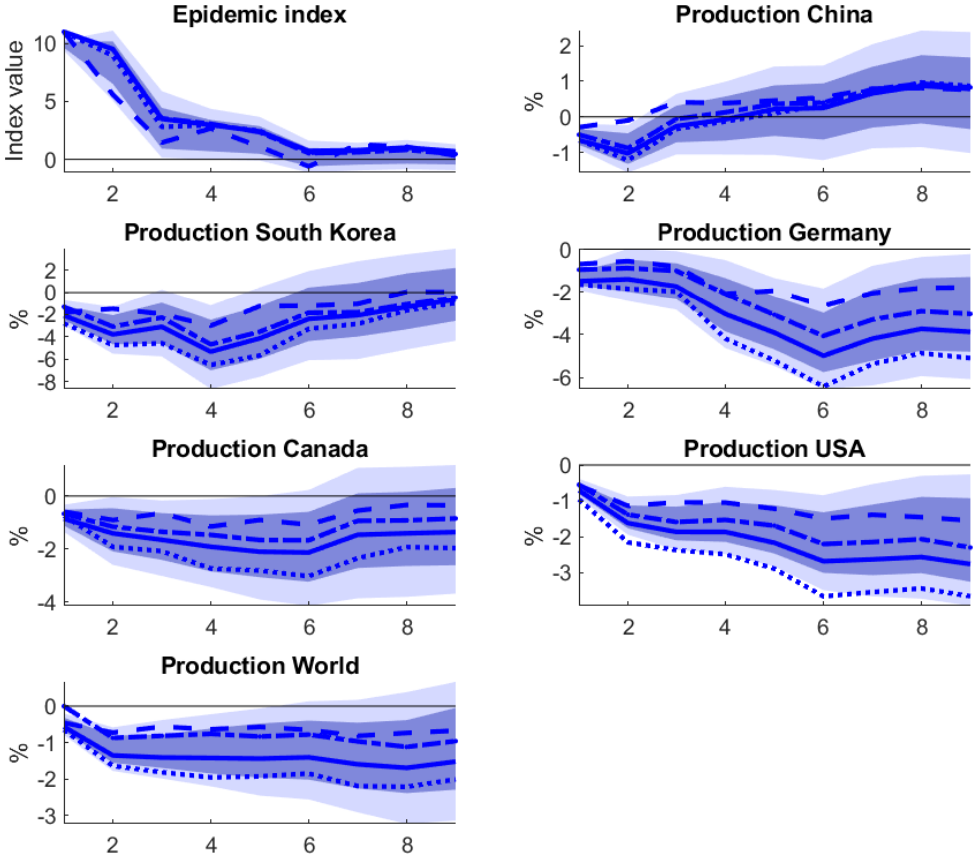

Figure 9: Responses to epidemic shock when addressing reverse causality concerns. *Notes:* The figure shows the responses of the endogenous variables to an epidemic shock of size 11 over a horizon of 9 months based on alternative SVAR(6). The solid lines and the shaded areas refer to the point estimates and the 68% and 90% confidence bands of the baseline specification, respectively, for comparison. The dashed lines refer to a model using an epidemic index based on the counting of keywords during the first week of each month only, disregarding weeks 2-4. The dotted lines refer to a model using a cleaned version of the epidemic index which controls for the joint occurrence of the disease words and 'economic/economy'. The dash-dotted lines refer to a model ordering world industrial production before the epidemic index.

production. The dash-dotted lines show that the results are qualitatively similar to the baseline findings. By construction, world production does not respond upon impact but then falls gradually. The trough response is halved relative, while the country-specific responses are only mildly affected. All in all, we conclude from the three tests that the main results are unlikely to be materially affected by reverse causality.

We also perform three tests to investigate the effect of potentially omitted variables. First, we compute each time the responses of the baseline variables when including one of the additional variables shown in Figures 4-7. These are 48 different models. Figures A.2-A.4 show that the main results hold. With a handful of exceptions, all the point estimates are covered by the confidence bands of the baseline model.

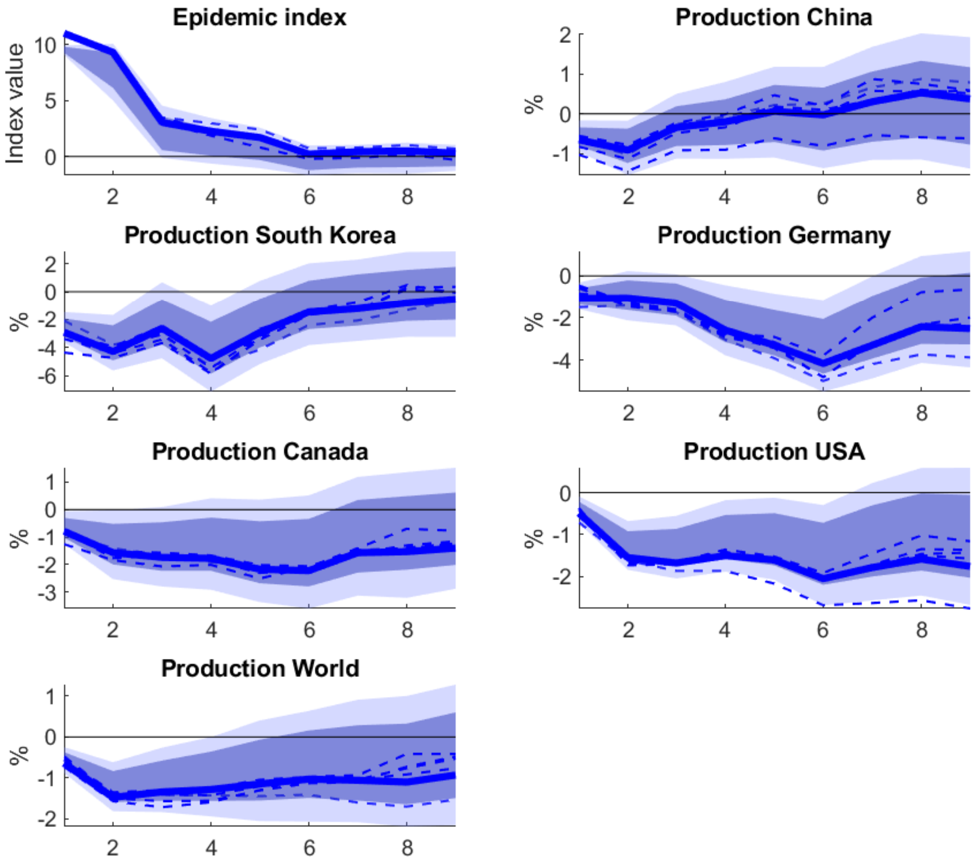

Figure 10: Responses to epidemic shock in factor-augmented models. *Notes:* The figure shows the responses of the endogenous variables to an epidemic shock of size 11 over a horizon of 9 months based on alternative SVAR(6) models. The solid lines and the shaded areas refer to the point estimates and the 68% and 90% confidence bands of the baseline specification, respectively, for comparison. The dashed lines refer to factor-augmented models adding one up to five factors jointly to the model. The factors are the principal components of 45 global and country-specific financial and real variables. They jointly account for 99.7% of the variation in these series.

Second, we condense the information in the additional variables into a few factors and estimate a factor-augmented structural VAR (Bernanke et al., 2005) that includes the available information jointly. We extract the first five principal components.[6] They explain jointly 99.7% of the variation in the data, with the first factor alone accounting for three fourth. Figure 10 replicates the baseline estimates for comparison (solid lines and shaded areas). The dashed lines refer to five alternative models that include one up to five factors, ordered last. Except for one model and the U.S., the responses for the

[6]We disregard the number of hotel accommodations for Hong Kong and the ifo service expectation and current situation surveys for Germany, which all start relatively late such that including them would shorten the common sample period substantially. This leaves us with 45 monthly data series, including both real and financial variables.

augmented models all lie within the bands of the baseline and the differences to the point estimates of the latter are small.

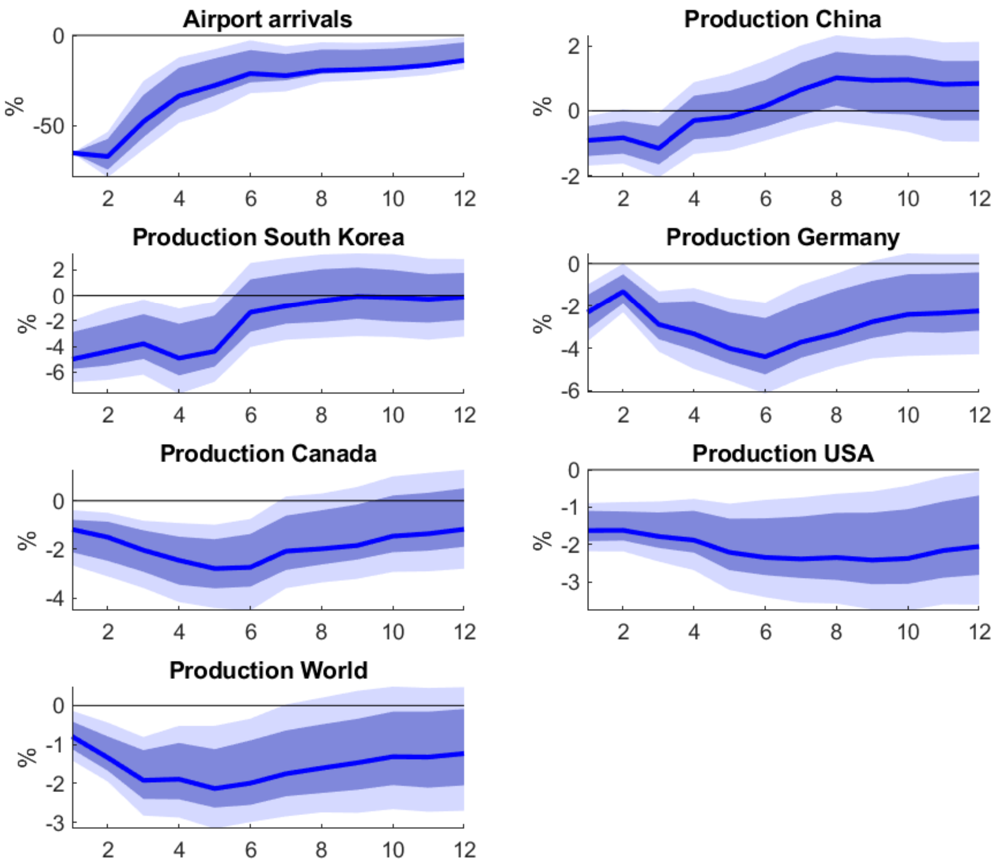

Figure 11: Responses to an epidemic shock in a Proxy-SVAR. *Notes:* The figure shows the responses of the endogenous variables to an epidemic shock over a horizon of 12 months based on a SVAR(6) identified through an external instrument. The shaded areas are 68% and 90% confidence bands. The instrument is the aggregate epidemic index and the indicator used to scale the latent shock the log of Hong Kong airport arrivals. The F-statistic for instrument relevance is 24.9.

Third, we construct an index for containment policy to separate the effects of epidemics from those of health interventions. We search for the words 'containment' (Eindämmung), 'vaccination' (Impfung), 'lockdown', 'mandatory mask wearing' (Maskenpflicht), 'restrictions on gatherings' (Personenobergrenze), 'quarantine' (Quarantäne), and 'school closure' (Schulschließung). We scale each series and then sum them. We include the index as endogenous variable into the model, replacing production in France. Figure A.5 shows a strong response of the containment index to the shock. However, the responses of the other variables are hardly affected compared to the baseline. Taken together, the three tests suggest that the results do not to suffer from omitted variables.

To determine whether the estimates are plagued by measurement error, we use the aggregate news index as an external instrument to identify epidemic shocks, following the methodology of Stock and Watson (2012) and Mertens and Ravn (2013). Online Appendix C.2 contains a description of the identification strategy. As indicator to scale the shock, we use the (log) airport arrivals at Hong Kong, the country with the highest mortality rate for SARS. The epidemic index is a strong instrument for the arrivals with an F-statistic of 24.9. We scale the latent epidemic shock to lower airport arrivals by 65%, which corresponds to the drop in that series from March to April 2003. Figure 11 shows that the main results hold. Although they are not directly comparable to the baseline estimates because of the alternative scaling of the shock, the effects tend to be larger and more significant, consistent with the idea that the instrumental approach addresses a potential attenuation bias.

Finally, we conduct technical robustness tests. We change the number of lags to $p = 3, 4, 5, 7, 8, 9$, include a linear trend or month dummies, winsorize the epidemic index at the 95th percentile, use local projections, and extend the model for the U.S. by employment and retail sales. Online Appendix C contains the results. The main findings hold.

6. Conclusions

We construct news-based viral disease indices and estimate the global economic impact of epidemics at the monthly frequency. We complement the analysis with country studies, an out-of-sample validation for Covid-19, and a historical study of the U.S. since 1920. The estimates provide new stylized facts. First, the economic damage of epidemic shocks is an order of magnitude larger than that of regular influenza, lasts for about one year, and is not recuperated. Second, the adverse impact of epidemics is roughly four times larger for epicenter countries than for countries indirectly affected. Third, epidemic shocks lead to a simultaneous fall of economic activity and consumer prices.

The results have several implications for public policy. They indicate that epidemics are costly tail events that are comparable to financial and political crises (Cerra and Saxena, 2008). As documented for those shocks, we find a high persistence and no significant evidence of overshooting. Hence, the hopes that the production losses in 2020/21 will be recuperated will likely be deceived. Moreover, the findings suggest that the negative demand effects of epidemic shocks are larger than the negative supply effects. These shocks are contractionary and deflationary. Together, these patterns indicate a role for expansionary fiscal policy in response to the Covid-19 crisis to reduce the permanent output losses. Finally, we hope that the indices will be fruitfully used by others in the future.

References

Adda, J., 2016. Economic activity and the spread of viral diseases: Evidence from high frequency data. The Quarterly Journal of Economics 131, 891–941.

Almond, D., 2006. Is the 1918 influenza pandemic over? long-term effects of in utero influenza exposure in the post-1940 us population. Journal of Political Economy 114, 672–712.

Baker, S.R., Bloom, N., Davis, S.J., 2016. Measuring economic policy uncertainty. The Quarterly Journal of Economics 131, 1593–1636.

Barro, R.J., Ursúa, J.F., Weng, J., 2020. The coronavirus and the great influenza pandemic: Lessons from the 'spanish flu' for the coronavirus's potential effects on mortality and economic activity. Technical Report. National Bureau of Economic Research.

Bernanke, B.S., Boivin, J., Eliasz, P., 2005. Measuring the effects of monetary policy: a factor-augmented vector autoregressive (favar) approach. The Quarterly Journal of Economics 120, 387–422.

Bisin, A., Gottardi, P., 2021. Efficient policy interventions in an epidemic. Journal of Public Economics 200, 104444.

Cerra, V., Saxena, S.C., 2008. Growth dynamics: the myth of economic recovery. American Economic Review 98, 439–57.

Clemens, J., Veuger, S., 2021. Politics and the distribution of federal funds: Evidence from federal legislation in response to covid-19. Journal of Public Economics 204.

Correia, S., Luck, S., Verner, E., 2020. Pandemics Depress the Economy, Public Health Interventions Do Not: Evidence from the 1918 Flu. Technical Report. mimemo.

Dixon, P.B., Lee, B., Muehlenbeck, T., Rimmer, M.T., Rose, A., Verikios, G., 2010. Effects on the us of an h1n1 epidemic: analysis with a quarterly cge model. Journal of Homeland Security and Emergency Management 7.

Eichenbaum, M., Rebelo, S., Trabandt, M., 2020. The macroeconomics of epidemics. mimemo, Freie Universität Berlin .

Fauci, A.S., Lane, H.C., Redfield, R.R., 2020. Covid-19 - navigating the uncharted. The New England Journal of Medicine 382, 1199–207.

Jordà, Ò., Singh, S.R., Taylor, A.M., 2020. Longer-run economic consequences of pandemics. Technical Report. National Bureau of Economic Research.

Karlsson, M., Nilsson, T., Pichler, S., 2014. The impact of the 1918 Spanish flu epidemic on economic performance in Sweden: An investigation into the consequences of an extraordinary mortality shock. Journal of Health Economics 36, 1–19.

Keogh-Brown, M.R., Wren-Lewis, S., Edmunds, W.J., Beutels, P., Smith, R.D., 2010. The possible macroeconomic impact on the UK of an influenza pandemic. Health Economics 19, 1345–1360.

Marinescu, I., Skandalis, D., Zhao, D., 2021. The impact of the federal pandemic unemployment compensation on job search and vacancy creation. Journal of Public Economics 200, 104471.

McKibbin, W., Sidorenko, A., et al., 2006. Global Macroeconomic Consequences of Pandemic Influenza. Technical Report. Lowy Institute for International Policy.

Mertens, K., Ravn, M.O., 2013. The dynamic effects of personal and corporate income tax changes in the united states. American Economic Review 103, 1212–1247.

Mitman, K., Rabinovich, S., 2021. Whether, when and how to extend unemployment benefits: Theory and application to covid-19. Journal of Public Economics , 104447–104447.

Pritchett, L., Summers, L.H., 1996. Wealthier is healthier. The Journal of Human Resources 31, 841–868.

Ramey, V.A., 2011. Identifying government spending shocks: It's all in the timing. The Quarterly Journal of Economics 126, 1–50.

Ruhm, C.J., 2000. Are recessions good for your health? The Quarterly Journal of Economics 115, 617–650.

Stock, J.H., Watson, M.W., 1999. Chapter 1 business cycle fluctuations in us macroeconomic time series, Elsevier. volume 1 of *Handbook of Macroeconomics*, pp. 3 – 64.

Stock, J.H., Watson, M.W., 2012. Disentangling the channels of the 2007-09 recession. Brookings Papers on Economic Activity , 120–157.

Verikios, G., Sullivan, M., Stojanovski, P., Giesecke, J., Woo, G., 2016. Assessing regional risks from pandemic influenza: a scenario analysis. The World Economy 39, 1225–1255.

Online appendix to ‘Viral shocks to the world economy’

Konstantin A. Kholodilin, DIW Berlin (German Institute for Economic Research), Mohrenstraße 58, 10117, Berlin, Germany; and National Research University Higher School of Economics, Kantemirovskaya ul., 3, 194100, St. Petersburg, Russia; kkholodolin@diw.de.

Malte Rieth, corresponding author, Martin-Luther-Universität Halle-Wittenberg, Universitätsring 3, 06108 Halle, Germany; and DIW Berlin; mrieth@diw.de.

December 8, 2021

Appendix A. Data definitions and sources

The news data are from the Genios and The New York Times archives. All other data are downloaded through Macrobond. Except for SARS-CoV-2, all data refer to the monthly frequency. We use seasonally adjusted real data, where available. Otherwise we make these adjustments ourselves. When constructing world aggregates ourselves, we confine the set of countries to strike a reasonable balance between having a sufficient number of countries and time-series observations to avoid changing compositions.

Variable	**Definition and transformation**
News-based disease indices	We measure the occurrence and intensity of viral diseases through news-based indices. To construct them, we use automated text analysis of two online media archives. The first one is the database Genios (www.genios.de). It includes about 2200 high-quality German-speaking media between 1990M1-2021M10 with the total number of documents exceeding 500 million. The second source is the archive of the New York Times (NYT), which spans the period 1910M1-2021M10 (www.nytimes.com/search?). For the estimation, we start the sample in 1923M1 to eliminate the loud noise in the influenza index at the very beginning of the sample (Figure 1). We search for the following five keywords that capture pandemics or internationally important epidemics during the sample period: ‘SARS’, ‘swine flu’ (Schweinegrippe), ‘MERS’, ‘Coronavirus’, and ‘Influenza’ (Grippe). We count the monthly occurrences of these keywords. Given that the amount of texts collected in the databases changes over time, we normalize the plain counts. We divide them by the number of occurrences of the word ‘der’, which is the most widely used word in German language, in the case of Genios, and by the occurrences of the word ‘new’ for the NYT archive, since the word ‘the’ is not searchable in that database. The aggregate news-based index for viral diseases summs of the the subindices for SARS, coronavirus and swine flu within months.
Deaths Germany	Total fatalities without still-born. Source: Federal Bureau of Statistics, Germany
Excess deaths Germany	Deviations of total fatalities from smoothed total fatalities, obtained through fitted third-order polynomial within month.
Deaths U.S.	Monthly and 12 month-ending number of live births, deaths and infant deaths. Source: Centers for Disease Control and Prevention, United Nations.
Excess deaths U.S.	Deviations of total fatalities from smoothed total fatalities, obtained through fitted second-order polynomial within month.
Infected and deaths SARS-CoV-2	Global number of infected persons with SARS-CoV-2 and deaths due to Covid-19. Sample: 2020M1D1-2020M4D6, excluding weekends. Source: WHO.
Global economic variables	
Industrial production	World Bank, Global Economic Monitor, Industrial Production, Total, Constant Prices, SA, USD, logarithm. Countries: world, China, South Korea, Germany, Canada, USA, Austria, Belgium, France, Luxembourg
Airport arrivals foreigners	Italy, International Arrivals, Holiday and Other Short-Stay Accommodation, Camping Grounds, Recreational Vehicle Parks and Trailer Parks, Foreign Countries, SA, Germany, Arrivals, Total, Foreigners [sa. X-11 ARIMA], Canada, CANSIM, Number of International Travellers Entering or Returning to Canada by Type of Transport, Canada, Total International Travellers [sa. X-11 ARIMA], Japan, Arrivals, Total Foreign Visitors [sa. X-11 ARIMA]. Logarithm of monthly sum.

Airport arrivals residents	Germany, Arrivals, Total, Residents [sa. X-11 ARIMA], Italy, International Arrivals, Holiday and Other Short-Stay Accommodation, Camping Grounds, Recreational Vehicle Parks and Trailer Parks, Italy, SA, New Zealand, Arrivals, By Type, NZ-Resident Travellers, Actual Counts [sa. X-11 ARIMA]. Logarithm of monthly sum.
Imports emerging markets	Foreign Trade, CPB World Trade Monitor, SA, Index, Emerging Markets, Import, Volume. Logarithm
Exports emerging markets	Foreign Trade, CPB World Trade Monitor, SA, Index, Emerging Markets, Export, Volume. Logarithm
Exports advanced countries	Foreign Trade, CPB World Trade Monitor, SA, Index, Advanced Economies, Export, Volume. Logarithm
Imports advanced countries	Foreign Trade, CPB World Trade Monitor, SA, Index, Advanced Economies, Import, Volume. Logarithm
Total trade world	Foreign Trade, CPB World Trade Monitor, SA, Index, World, Total, Volume. Logarithm
Employment world	Germany, Employment, Total, Domestic Concept, SA (X13 JDemetra+); United States, Employment, National, 16 Years and Over, SA; Austria, Employment, Employed Persons, Total, Persons in Dependent Employment [sa. X-11 ARIMA]; Japan, Employment, Employed Persons, Total, National, Males and Females, SA; United Kingdom, Employment, Aged 16-64, SA; Canada, Employment, Women and Men, 15 Years and Over, SA; Taiwan, Employment, Total [sa. X-11 ARIMA] ; Australia, Employment, Total, SA; South Korea, Labor Force Statistics, Economically Active Persons, Employed Persons [sa. X-11 ARIMA]; Hong Kong, Employment, Total [sa. X-11 ARIMA]. Logarithm of monthly sum.
Consumer prices world	World Bank, Global Economic Monitor, Prices, Consumer Price Index, SA, Index. Logarithm
Retail sales world	Austria, Domestic Trade, Retail Trade, Turnover, Total, Excluding Trade in Motor Vehicles, Constant Prices, Index; Germany, Domestic Trade, Retail Trade, Turnover, Total, Excluding Vehicle Trade, Calendar Adjusted (X13 JDemetra+), Constant Prices, SA (X13 JDemetra+), Index; Australia, Domestic Trade, Retail Trade, By Industry, Total, Current Prices, SA, AUD; Japan, Domestic Trade, Retail Trade, Total, JPY [sa. X-11 ARIMA] ; Hong Kong, Domestic Trade, Retail Trade, Total Sales, Value, HKD [sa. X-11 ARIMA] ; Singapore, Domestic Trade, Wholesale and Retail Trade, Retail Sale, Total, Constant Prices, SA, Index; Mexico, Domestic Trade, Retail Trade, Total, Constant Prices, SA, Index; Sweden, Domestic Trade, Retail Trade, Total except Fuel, SA, Index; Canada, Domestic Trade, Retail Trade, Total, SA, CAD; Portugal, Domestic Trade, Retail Trade, Total, Excluding Fuel, Index [sa. X-11 ARIMA] ; United States, Domestic Trade, Retail Trade, Retail Sales, Total, Calendar Adjusted, SA, USD. Aggregation of individual growth rates using fixed-GDP weights as of 2010. Logarithm of aggregate index.
Equity prices world	Equity Indices, MSCI, Large Cap, Index, Total Return, Local Currency; World. Logarithm
Federal funds rate	United States, Policy Rates, Effective Rates, Federal Funds Effective Rate

Economic variables for Hong Kong

Airport arrivals	Arrivals, Total [sa. X-11 ARIMA]. Logarithm
Airport departures	Departures, By Border Checkpoint, Total [sa. X-11 ARIMA]. Logarithm
Accommodation occupancy rate	Accommodation, Occupancy, Hotels, Room Occupancy Rate, All Hotels [sa. X-11 ARIMA]
Exports	Foreign Trade, Total Exports, Quantum, Index [sa. X-11 ARIMA]. Logarithm
Imports	Foreign Trade, Imports, Quantum, Index [sa. X-11 ARIMA]. Logarithm
Employment	IMF IFS, Real Sector, Labor, Employment, Persons [sa. X-11 ARIMA]. Logarithm
Unemployment rate	IMF IFS, Real Sector, Labor, Unemployment Rate
Retail sales total	Hong Kong, Domestic Trade, Retail Trade, Total Sales, Value, HKD [sa. X-11 ARIMA]. Logarithm
Consumer prices	Consumer Price Index, Index. Logarithm
Government revenues	Government Fiscal Operations, Revenues, Provisional, HKD [sa. X-11 ARIMA]. Logarithm
Government expenditures	Government Fiscal Operations, Expenditures, Provisional, HKD [sa. X-11 ARIMA]. Logarithm
Policy rate	Policy Rates, Central Bank Policy Rate, End of Period
Sales jewellery, watches, gifts	Domestic Trade, Retail Trade, Jewellery, Watches and Clocks, and Valuable Gifts, Total, Per CPI, HKD [CPI index Oct. 2014 - Sep. 2015 = 100, sa. X-11 ARIMA]. Logarithm
Sales clothing and footwear	Domestic Trade, Retail Trade, Clothing, Footwear and Allied Products, Total, Per CPI, HKD [CPI index Oct. 2014 - Sep. 2015 = 100, sa. X-11 ARIMA]. Logarithm
Sales department stores	Domestic Trade, Retail Trade, Department Stores, Total, Per CPI, HKD [CPI index Oct. 2014 - Sep. 2015 = 100, sa. X-11 ARIMA]. Logarithm
Sales other consumer goods	Domestic Trade, Retail Trade, Other Consumer Goods, Total, Per CPI, HKD [CPI index Oct. 2014 - Sep. 2015 = 100, sa. X-11 ARIMA]. Logarithm
Sales consumer durables	Domestic Trade, Retail Trade, Consumer Durable Goods, Total, Per CPI, HKD [CPI index Oct. 2014 - Sep. 2015 = 100, sa. X-11 ARIMA]. Logarithm
Sales supermarkets	Domestic Trade, Retail Trade, Supermarkets, Total, Per CPI, HKD [CPI index Oct. 2014 - Sep. 2015 = 100, sa. X-11 ARIMA]. Logarithm

Economic variables for Germany

Imports	Foreign Trade, Total, Calendar Adjusted (X13 JDemetra+), SA (X13 JDemetra+), EUR, Imports. Logarithm
Foreign orders	Foreign orders, real, SA. Logarithm
Exports	Foreign Trade, Total, Calendar Adjusted (X13 JDemetra+), SA (X13 JDemetra+), EUR, Exports. Logarithm

Capacity utilization	Capacity Utilization, Manufacturing, Total, SA (X-13 ARIMA)
Retail sales	Retail sales, real, SA. Logarithm
Domestic orders	Domestic orders, real, SA. Logarithm
ifo situation manufacturing	ifo situation manufacturing
ifo expectations manufacturing	ifo expectations manufacturing
ifo situation services	ifo situation services
ifo expectations services	ifo expectations services
Short-time work	Short-time work. Logarithm
Employment	Employment, total, SA. Logarithm
Production intermediate goods	Industrial Production, By Goods, Calendar Adjusted (X13 JDemetra+), Constant Prices, SA (X13 JDemetra+), Index, Production intermediate goods. Logarithm
Production capital goods	Industrial Production, By Goods, Calendar Adjusted (X13 JDemetra+), Constant Prices, SA (X13 JDemetra+), Index, Production capital goods. Logarithm
Production non-durables	Industrial Production, By Goods, Calendar Adjusted (X13 JDemetra+), Constant Prices, SA (X13 JDemetra+), Index, Production non-durables goods. Logarithm
Production consumer goods	Industrial Production, By Goods, Calendar Adjusted (X13 JDemetra+), Constant Prices, SA (X13 JDemetra+), Index, Production consumer goods. Logarithm
Producer prices	Producer Price Index, Industrial Products, Total, Calendar Adjusted, SA, Index. Logarithm
Consumer prices	Consumer Price Index, Total, Calendar Adjusted, SA, Index. Logarithm
Economic variables for the United States	
Industrial production	Industrial Production, Constant Prices, SA, Index, Total. Logarithm
Employment	Employment, Payroll, SA, Nonfarm, Total. Logarithm
S&P 500	Equity Indices, S&P, 500, Index, Total Return, End of Period, USD. Logarithm
Consumer prices	Consumer Price Index, SA. Logarithm
Retail sales	OECD MEI, Sales, Retail Trade, SA, Total Retail Trade, Volume, Index. Logarithm
Commercial paper rate 3m	Commercial Paper Rates, Rates, AA Nonfinancial, 3 Month, Yield
Oil and gas production	EIA, Oil and Gas, Crude Oil, Production, Total, Barrels. Logarithm

Appendix B. Details of epidemics

This section provides details on the epidemics in the base sample and additional estimation results.

SARS. The outbreak of this disease caused by the SARS coronavirus (SARS-CoV) originated in Foshan, China in November 2002 and lasted until May 2004. Between November 1, 2002, and July 31, 2003, 8,096 people were infected and at least 774 of them died worldwide. The most affected regions were Mainland China with 5,327 cases and Hong Kong with 1,755 cases.[7] In Figure 1, the index peaks in April 2003; that is, half a year after the outbreak of the disease. Another, much smaller peak, in February 2020, is related to the comparison of Covid-19 to SARS in the media.

Swine flu. The swine flu was an influenza pandemic caused by the H1N1/09 virus. It was first recognized in California and Texas in March 2009. The pandemic lasted from April 2009 to July 2010.[8] In terms of the number of confirmed deaths, the pandemic affected mostly the USA, Brazil, India, and Mexico. By July 2009, there were 94,512 confirmed cases, including 429 deaths.[9]

[7] The WHO summary of probable SARS cases with onset of illness is from November 1, 2002, to July 31, 2003.

[8] European Centre for Disease Prevention and Control, Timeline on the pandemic (H1N1) 2009; https://www.ecdc.europa.eu/en/seasonal-influenza/2009-influenza-h1n1-timeiline.

[9] World Health Organization, Human infection with pandemic (H1N1) 2009 virus: updated interim WHO guidance on global surveillance, p. 7, 2009.

MERS. The 'Middle East respiratory syndrome' is caused by the betacoronavirus (MERS-CoV). As the name suggests, the epidemic originated in the Middle East, specifically, Saudi Arabia, where approximately 80% of cases were identified. The first laboratory-confirmed case was reported in Saudi Arabia in April 2012. The first large outbreak of MERS started in March 2014 and ended in July 2015. Since then, and through November 2019, there were 2,494 laboratory-confirmed cases, including 858 deaths.[10]

COVID-19. The disease is caused by the SARS-CoV-2 coronavirus. Its outbreak was identified in December 2019 in the city of Wuhan, Hubei province, China. By March 2020, there are about 1 million confirmed cases, including more than 50,000 deaths.[11] The corresponding index started to rise in January 2020. Already at this early stage, the intensity of media coverage is close to that of SARS or swine flu at their peaks, and much stronger than for MERS.

Influenza. The figure also contains indices of influenza for Germany based on the Genios database and for the U.S. based on the NYT archive. For the overlapping period, the correlation between the indicators is 0.64. Both have recurrent medium-sized peaks, typically during the winter. For Germany, milder courses of the disease are associated with peaks of 2-3, while the largest non-epidemic spike is at 7. For the U.S., typical seasonal peaks are close to, but below, 10 and the largest non-epidemic spike is at 20. In addition, both indices have several outliers. They capture epidemic or pandemic influenza. For Germany, they correspond to the global avian flu (H5N1 virus) and swine flu outbreaks in 2005 and 2009, respectively. For the U.S., where the sample starts much earlier, the large spikes correspond to the domestic epidemic influenza of 1929, to the 1957/58 pandemic Asian flu (H2N2 virus), to the 1968/69 pandemic Hong Kong flu (H3N2 virus), as well as to the avian and swine flus that are also present in the German series.

Appendix C. Sensitivity analysis

This section provides further details of the sensitivity analysis.

Appendix C.1. Addressing potential sources of endogeneity

To investigate the effect of potentially omitted variables, we compute the responses of the baseline variables in the model for the world economy (Figure 3) when replacing the production in France by the additional variables shown in Figures 4-6. The main results hold (see Figures A.2-A.4). This finding reflects that the baseline model already contains output of the world as a whole and of main economies individually such that the potential for omitting important variables is low in the first place.

[10] World Health Organization, MERS situation update, November 2019.
[11] World Health Organization, Coronavirus disease 2019 (COVID-19), Situation Report - 45, 2020.

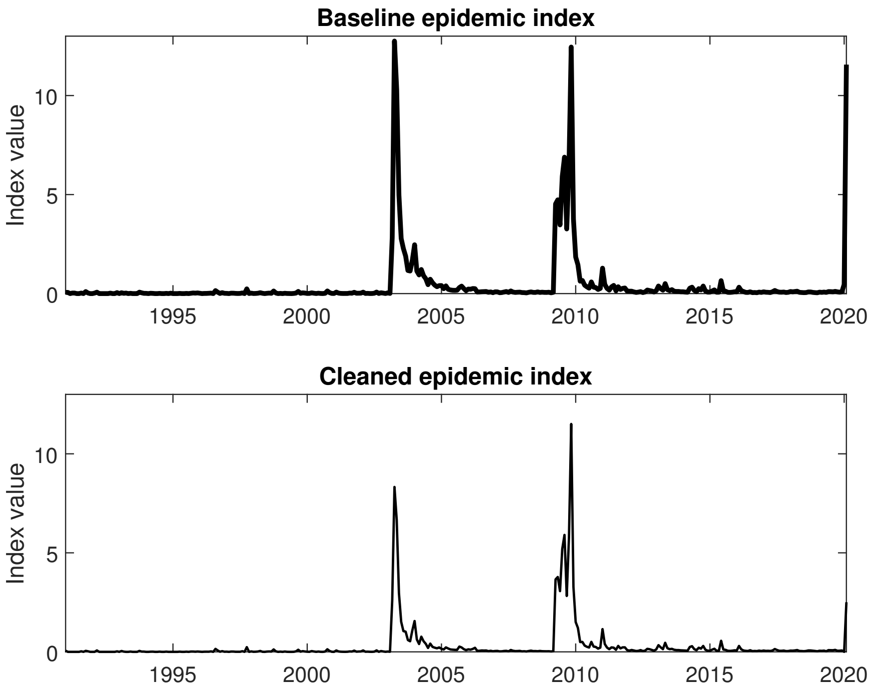

Figure A.1: Comparison of baseline and cleaned epidemic index. *Notes:* The figure shows the baseline aggregate epidemic index (upper panel) and a cleaned version thereof (lower panel). The latter is obtained by counting the joint occurrence of the words, say, 'coronavirus' and 'economic/economy' for each disease subindex and then subtracting the normalized counts of the joint hits for the subindices from the normalized aggregate index.

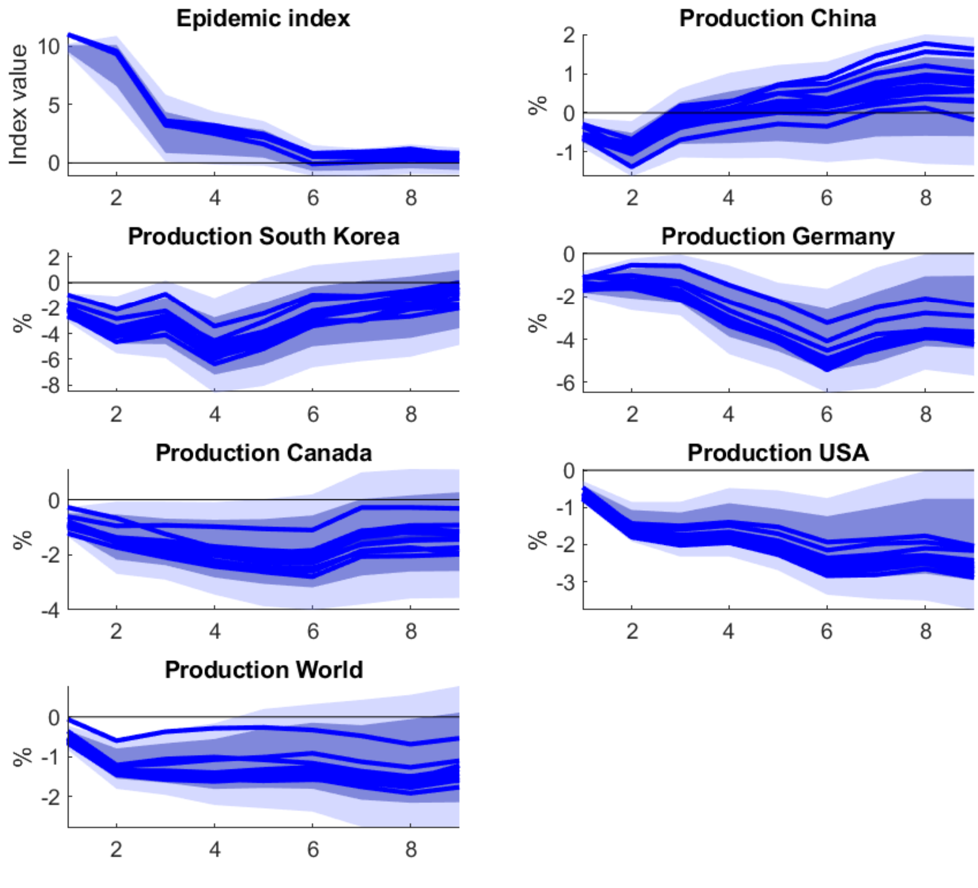

Figure A.2: Robustness to adding further global variables. *Notes:* The figure shows the responses of the endogenous variables to an epidemic shock of size 11 over a horizon of 9 months based on a SVAR(6). The 68% and 90% confidence bands refer to the baseline specification. The solid lines show the point estimates of the baseline variables in the augmented models where the variables shown in Figure 4 are added one at a time.

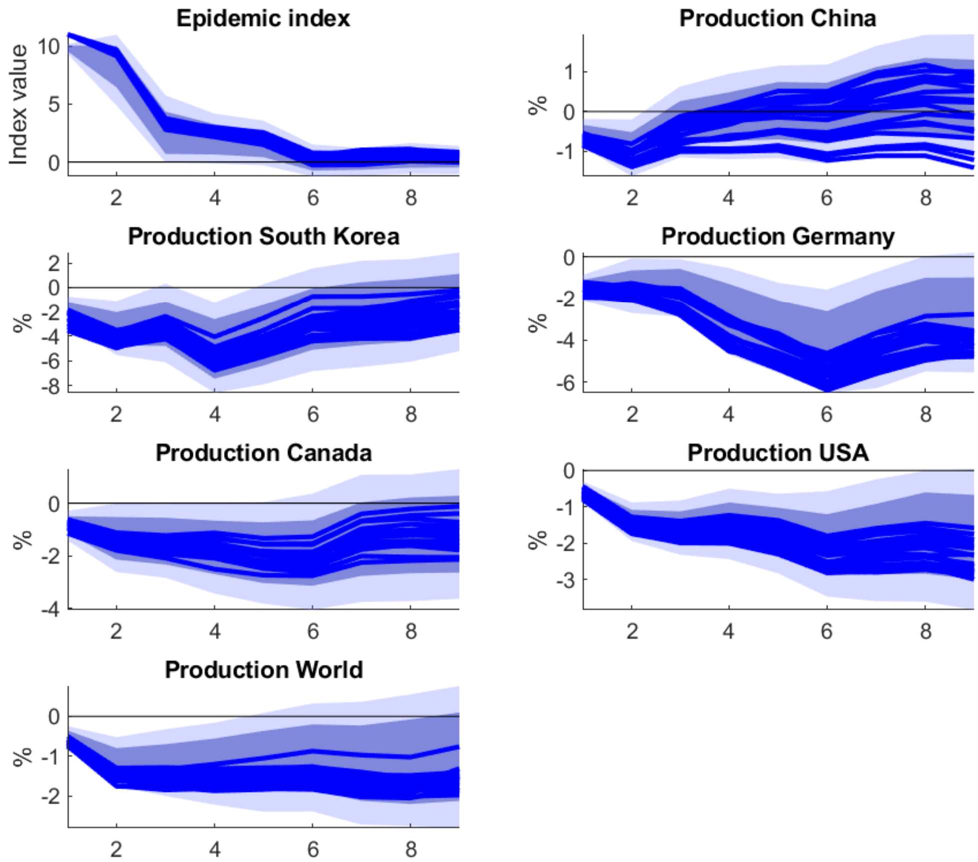

Figure A.3: Robustness to adding variables for Hong Kong. *Notes:* The figure shows the responses of the endogenous variables to an epidemic shock of size 11 over a horizon of 9 months based on a SVAR(6). The 68% and 90% confidence bands refer to the baseline specification. The solid lines show the point estimates of the baseline variables in the augmented models where the variables shown in Figure 6 are added one at a time.

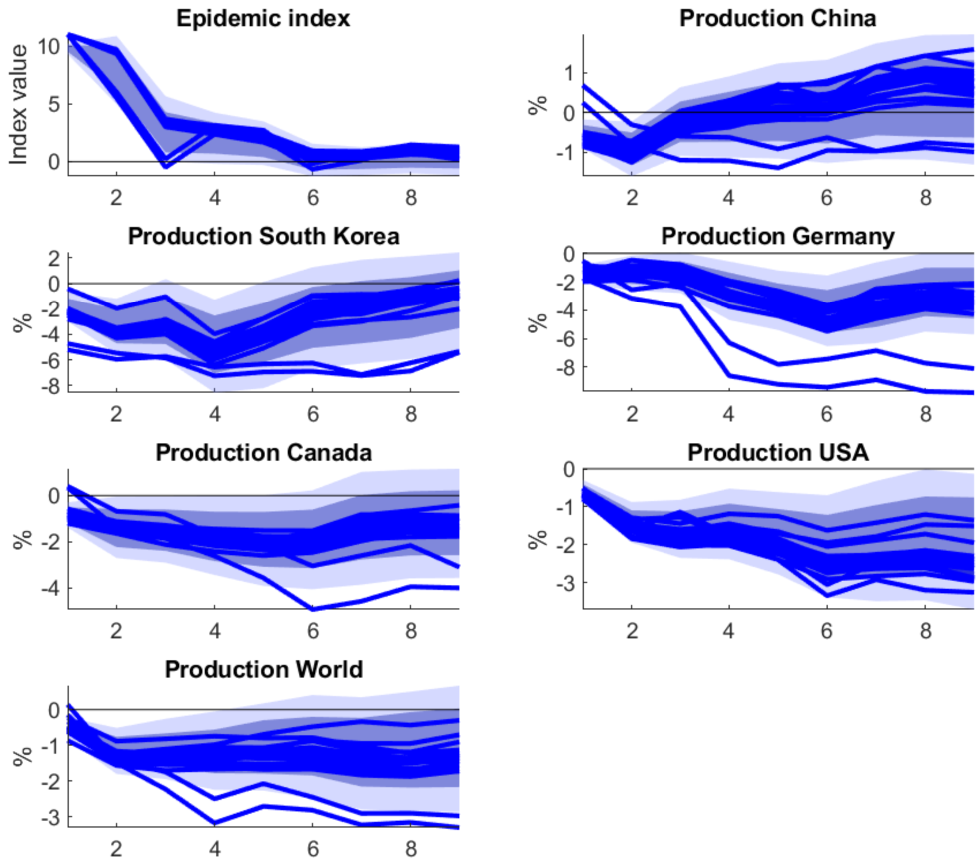

Figure A.4: Robustness to adding variables for Germany. *Notes:* The figure shows the responses of the endogenous variables to an epidemic shock of size 11 over a horizon of 9 months based on a SVAR(6). The 68% and 90% confidence bands refer to the baseline specification. The solid lines show the point estimates of the baseline variables in the augmented models where the variables shown in Figure 7 are added one at a time.

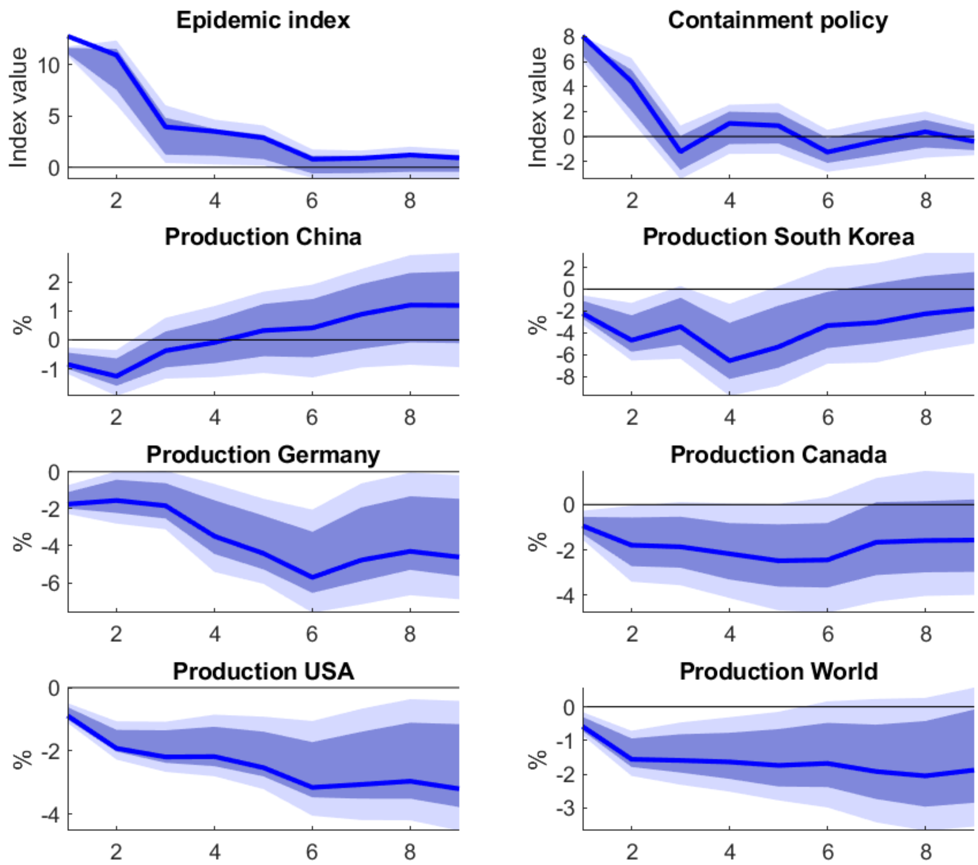

Figure A.5: Robustness to adding containment policy index. *Notes:* The figure shows the responses of the endogenous variables to an epidemic shock of size 11 over a horizon of 9 months based on a SVAR(6), together with 68% and 90% confidence bands. The model includes an index of containment policy index.

To determine if the estimates are plagued by measurement error in the aggregate news-based disease index, we follow the external instrument approach for SVARs developed in Stock and Watson (2012) and Mertens and Ravn (2013). We assume that the index is an instrumental variable s_t that is correlated with the latent epidemic shock of interest, but uncorrelated with other structural shocks and hence fulfills

$$\mathbb{E}[s_t \varepsilon_{1t}] = \phi \neq 0 \tag{A.1}$$

$$\mathbb{E}[s_t \varepsilon_{jt}] = 0 \quad \forall\, j = 2, \ldots, K, \tag{A.2}$$

where ϕ is an unknown correlation between the instrument s_t and the structural shock of interest ε_{1t}. The latter is ordered first without loss of generality. In the literature, (A.1) is usually called the relevance condition and assumption (A.2) the exogeneity condition. A valid instrument satisfies both (A.1) and (A.2). It allows for recovering ε_{1t} and, hence, the corresponding response vector from the reduced form residuals. Using $B = [b_1, B^*]$, where b_1 is the response vector corresponding to ε_{1t} and B^* contains the responses of the remaining shocks, yields

$$u_t = b_1 \varepsilon_{1t} + B^* \varepsilon_t^*. \tag{A.3}$$

Substituting (A.3) into $\mathbb{E}(s_t u_t)$, while using (A.1) and (A.2), allows for uncovering the (relative) impact of the structural shock of interest on every variable in the system, that is, the jth element of b_1. By using the sample moments $\hat{\mathbb{E}}(u_t s_t)$, the instrument s_t implies the following $k-1$ identifying restrictions

$$b_1 = b_{11} \left(1, \frac{\hat{\mathbb{E}}(u_{2t} s_t)}{\hat{\mathbb{E}}(u_{1t} s_t)}, \ldots, \frac{\hat{\mathbb{E}}(u_{Kt} s_t)}{\hat{\mathbb{E}}(u_{1t} s_t)}\right)', \tag{A.4}$$

posing identification of shock ε_{1t} up to the scaling factor b_{11}. To scale the shock, we need an indicator variable that enters the SVAR. We use the log of the number of airport passenger arrivals in Hong Kong, which was the economy most affected in the world in terms of relative SARS cases. The epidemic index is a strong instrument. It has a F-statistic of 24.9.

Appendix C.2. Further sensitivity tests

We also conduct several more technical robustness tests. Instead of using 6 lags, we change the lag length to 3, 4, 5, 7, 8, and 9, respectively (Figure A.7). Next, we incorporate a linear trend into the model (solid lines and shaded area in Figure A.8), month dummies (dashed lines in Figure A.8), or we winsorize the epidemic index at the 95th percentile (dotted lines in Figure A.8). All-in-all, the results hold. Finally, we extend the model for the U.S. by including the logarithm of employment and of retail sales. This shortens the sample by half to start in 1955M1. The drop in prices is now more significant. Employment falls and retail sales increase, potentially reflecting stockpiling by consumers, but the responses of the two additional variables are only borderline significant.

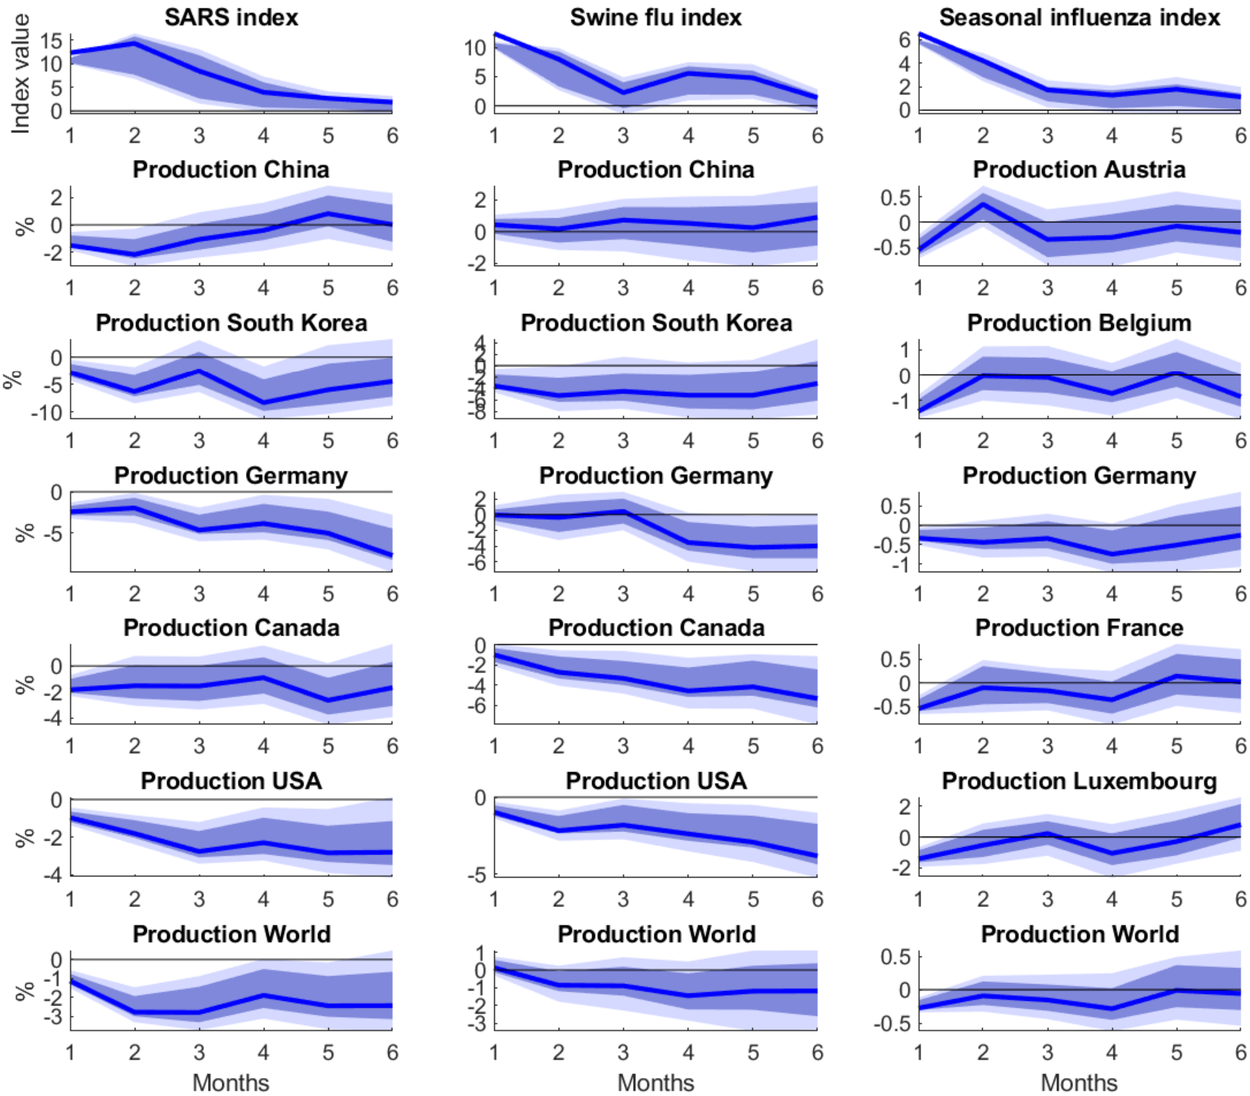

Figure A.6: SARS, swine flu, and regular influenza shocks when controlling for seasonality. *Notes:* The figure shows the dynamic impact of a SARS shock (column 1), a swine flu shock (column 2), and a regular influenza shock (column 3), obtained from three SVAR models, on production in single countries and globally over 6 months, along with 68% and 90% bootstrapped confidence bands, when controlling for month-of-the-year effects.

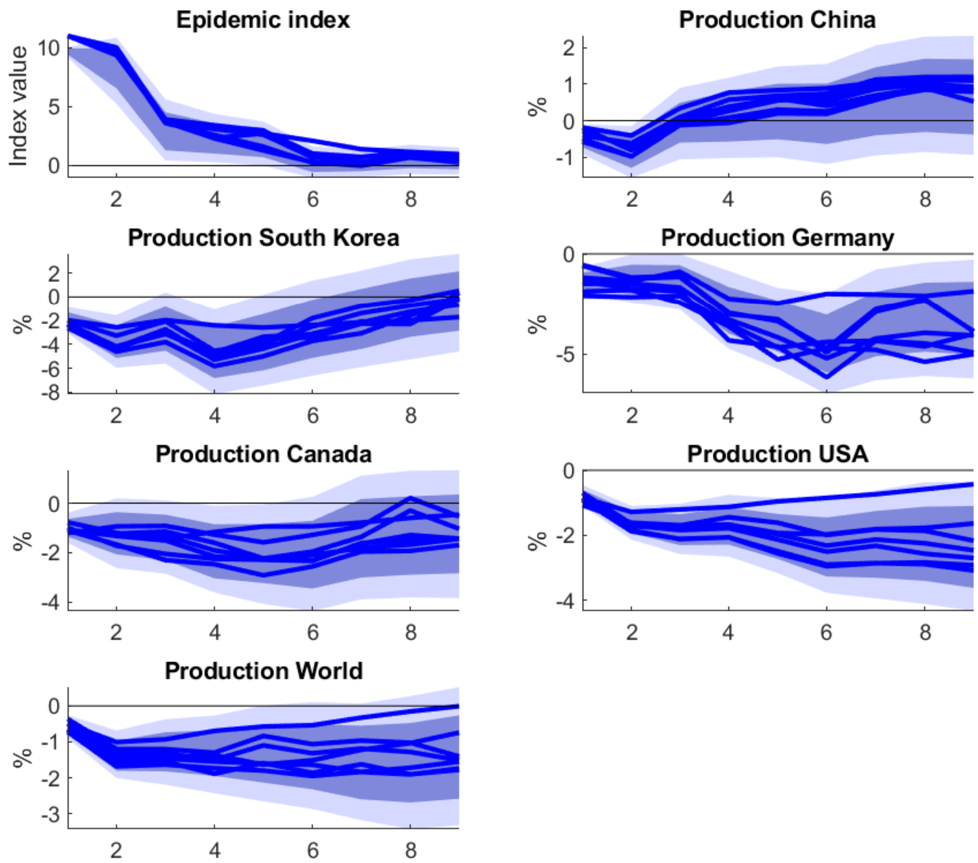

Figure A.7: Responses to an epidemic shock using alternative lag length. *Notes:* The figure shows the responses of the endogenous variables to an epidemic shock of size 11 over a horizon of 9 months based on a SVAR(p), with $p = 3, 4, 5, 6, 7, 8, 9$, respectively. The shaded areas are 68% and 90% confidence bands and refer to the baseline specification with $p = 6$.

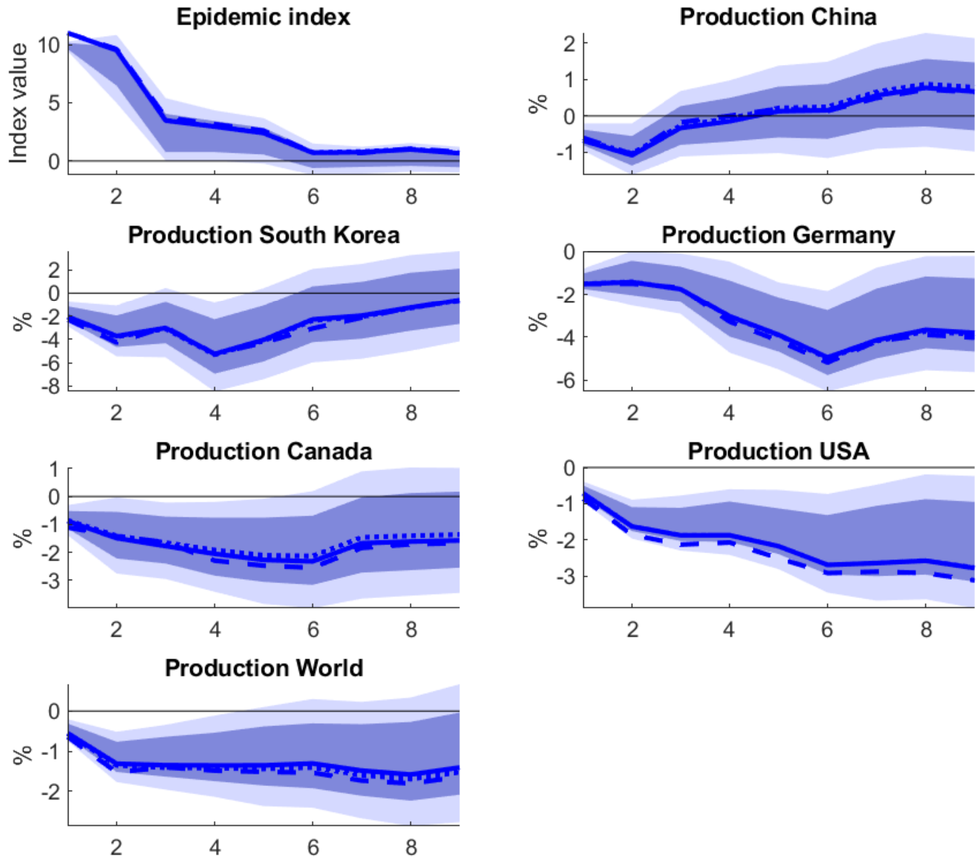

Figure A.8: Responses to an epidemic shock when including a trend, month dummies, or a winsorized epidemic index. *Notes:* The figure shows the responses of the endogenous variables to an epidemic shock of size 11 over a horizon of 9 months based on a SVAR(6). The solid lines and the shaded areas for 68% and 90% confidence bands refer to a model including a linear trend, the dashed lines to a model with month dummies, and the dotted lines to a model where we winsorize the epidemic indicator at the 95th percentile.

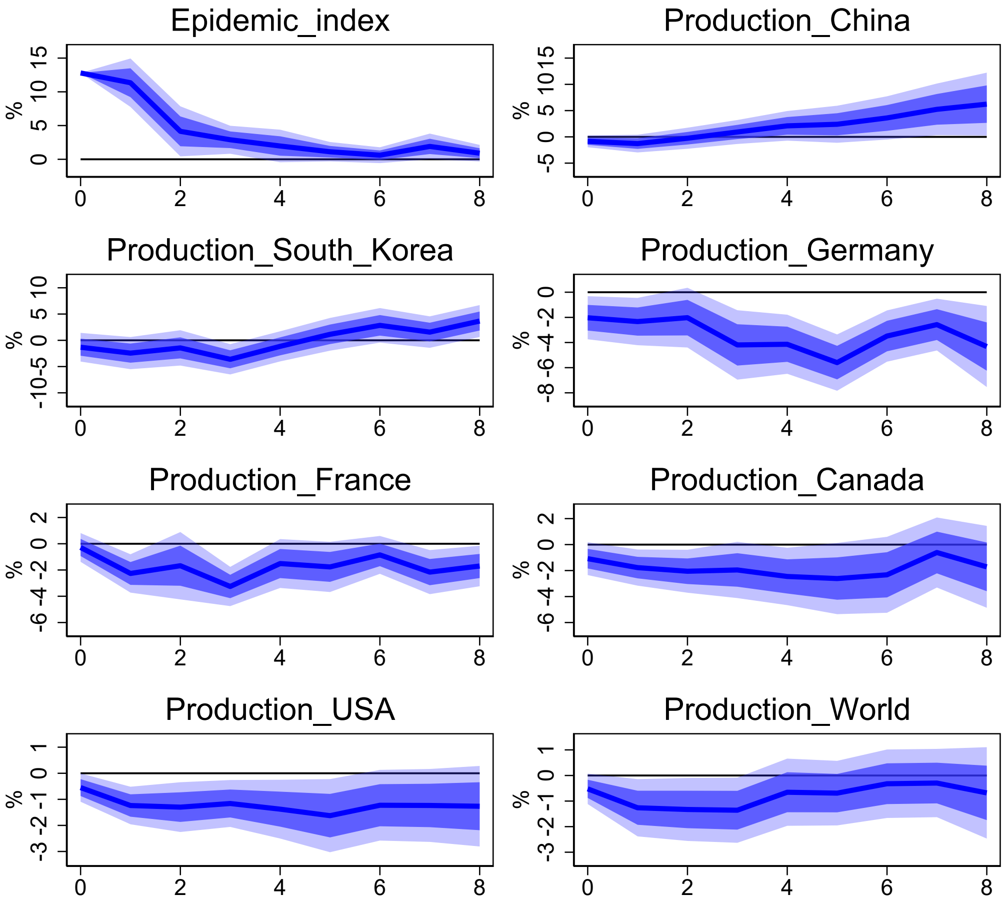

Figure A.9: Responses to an epidemic shock using local projections. *Notes:* The figure shows the responses of the endogenous variables to an increase in the epidemic index of size 12.8 over a horizon of 9 months based on local projections with six lags of the variables in the baseline global model as controls. The shaded areas are 68% and 90% hereroskedasticity and autocorrelation robust confidence bands with the number of lags equal to the response horizon plus one.

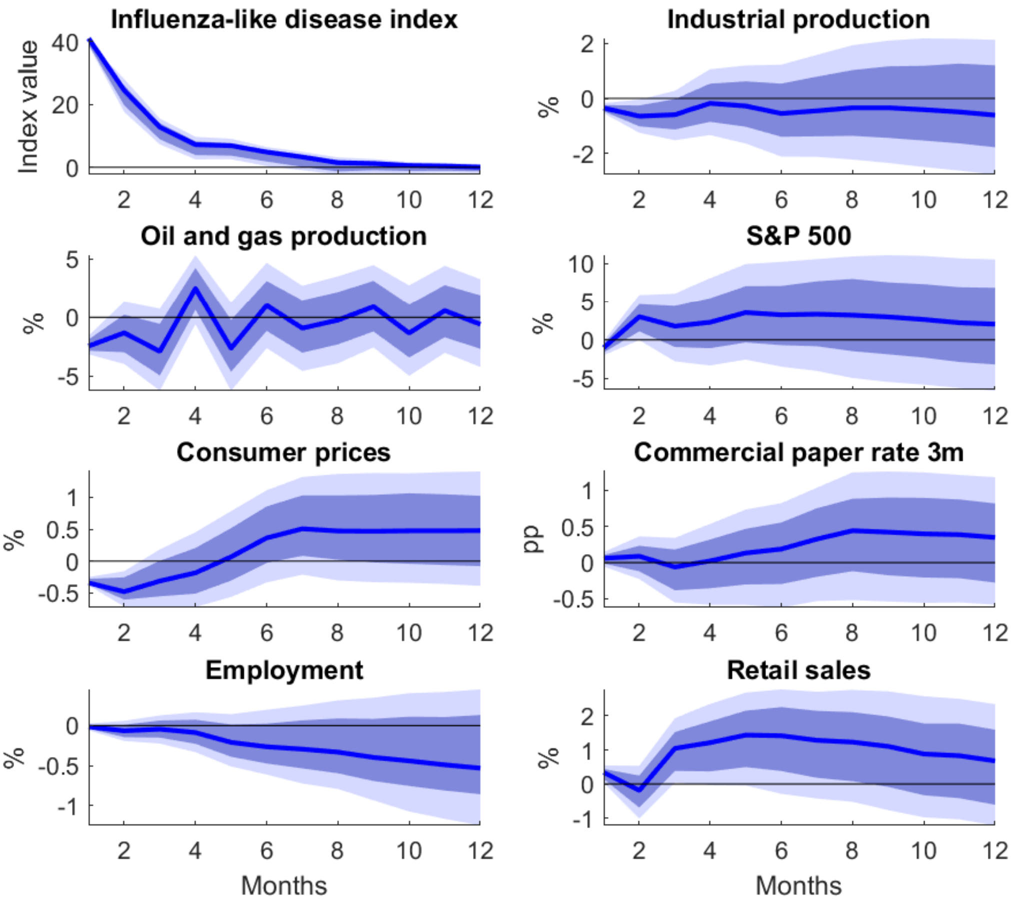

Figure A.10: Responses of U.S. variables to an influenza-like disease shock in extended model. *Notes:* The figure shows the responses of the endogenous variables to an influenza shock of size 41 over a horizon of 12 months based on a SVAR(6) for the U.S. The shaded areas indicate 68% and 90% confidence bands. The extended model includes additionally the logarithm of employment and of retail sales.