People Meet People: A Microlevel Approach to Predicting the Effect of Policies on the Spread of COVID-19

Janoś Gabler
Tobias Raabe
Klara Röhrl

I Z A Institute
of Labor Economics

Initiated by Deutsche Post Foundation

I Z A Institute
of Labor Economics
Initiated by Deutsche Post Foundation

Janoś Gabler
Bonn Graduate School of Economics and IZA

Tobias Raabe

Klara Röhrl
Bonn Graduate School of Economics

ABSTRACT

People Meet People: A Microlevel Approach to Predicting the Effect of Policies on the Spread of COVID-19

Governments worldwide are adopting nuanced policy measures to reduce the number of Covid-19 cases with minimal social and economic costs. Epidemiological models have a hard time predicting the effects of such fine grained policies. We propose a novel simulation-based model to address this shortcoming. We build on state-of-the-art agent-based simulation models but replace the way contacts between susceptible and infected people take place. Firstly, we allow for heterogeneity in the types of contacts (e.g. recurrent or random) and in the infectiousness of each contact type. Secondly, we strictly separate the number of contacts from the probabilities that a contact leads to an infection. The number of contacts changes with social distancing policies, the infection probabilities remain invariant. This allows us to model many types of fine grained policies that cannot easily be incorporated into other models. To validate our model, we show that it can accurately predict the effect of the German November lockdown even if no similar policy has been observed in the time series that were used to estimate the model parameters.

Corresponding author:
Janoś Gabler
IZA
Schaumburg-Lippe-Str. 5-9
D-53113 Bonn
Germany
E-mail: gabler@iza.org

1 Introduction

The first wave of the Covid-19 pandemic prompted strict lockdowns and restrictions across the world. As a result, many countries were able to reduce the number of positive cases and could relax restrictions imposed on their societies. At the same time, the social and economic costs were enormous. In the second wave, countries are trying to implement more nuanced policies. For example, Germany imposed a "lockdown-light" in November. Some businesses are closed and certain leisure activities are prohibited but schools and nurseries stay open.

Current epidemiological models have not been designed to predict the effect of such fine-grained policies. These models need to be extended for each new policy proposal and cannot be easily adapted to fast changing environments. This report describes a model that has been designed from the ground up to predict the effects of contact reducing policies in real time. At the time of this writing, it has the following features:

1. At the core of the model, people meet people based on a matching algorithm. We distinguish various types of contacts. Currently, these are households, leisure activities, schools, nurseries and several types of contacts at the workplace. Contact types can be random or recurrent and vary in frequency.

2. Policies can be implemented as shutting down contact types entirely or partially. The reduction of contacts can be random or systematic, e.g., to allow for essential workers.

3. Infection probabilities vary across contact types, but are invariant to policies which reduce contacts.

4. The model achieves a good fit on German infection and fatality rate data even if only the infection probabilities are fit to the data and the remaining parameters are calibrated from the medical literature and datasets on contact frequencies.

5. High quality Python code for the model is freely available on https://github.com/covid-19-impact-lab/sid, well documented and very flexible. We are actively looking for researchers who want to use our model for their projects.

After a brief literature review we describe our model in detail and validate its in-sample and out-of-sample fit on German data.

This report describes the model in an abstract way, but uses many realistic examples from a version that is specialized to Germany. It is important to note that this specialization is not baked into the model or the Python code. It is easy to adjust the model to other countries if data on the number of contacts and a dataset with background characteristics are available.

More details about the German model as well as applications to currently discussed German policies can be found in a companion paper (Dorn et al., 2020a).

2   Literature Review

We build on two strands of literature: Recent extensions of the epidemiological SEIR model and agent-based simulation models.

The traditional SEIR model is not fine-grained enough to model nuanced policies. This has motivated a large number of researchers to extend the standard model to allow for more heterogeneity and flexibility. Examples are Grimm et al., 2020, Donsimoni et al., 2020 and Acemoglu et al., 2020 who develop multi group SEIR models to analyze the effects of targeted lockdowns and Berger et al., 2020 who extend the SEIR model to analyze testing and conditional quarantines. For a more comprehensive review see Avery et al., 2020. Others have used the results of a standard SEIR model as input for economic models that estimate the cost of policies (e.g. Dorn et al., 2020b).

While the popularity of the SEIR model is mainly due to its simplicity, the extensions are quite complex. It is unlikely that there will be a SEIR model that combines all proposed extensions. Moreover, the extensions do not address other key issues: The main parameter of the SEIR model, the basic reproduction number (R_0), is not policy-invariant. It is a composite of the number of contacts each person has and the infection probability of the contacts. In fact, policy simulations are done by setting R_0 to a different value but it is hard to translate a real policy into the value of R_0 it will induce.

Another commonly used model class in epidemiology are agent-based simulation models. In these models individuals are simulated as moving particles. Infections take place when two particles come closer than a certain contact radius (e.g. Silva et al., 2020 and Cuevas, 2020). While the simulation approach makes it easy to incorporate heterogeneity in disease progression, it is hard to incorporate heterogeneity in meeting patterns. Moreover, policies are modeled as changes in the contact radius or momentum equation of the particles. The translation from real policies to corresponding model parameters is a hard task.

Hinch et al., 2020 is a recent extension of the prototypical agent-based simulation model that replaces moving particles by contact networks for households, work and random contacts. This model is similar in spirit to ours but focuses on contact tracing rather than social distancing policies.

The above assessment of epidemiological models is not meant as a critique. We are aware that those models were not designed to predict the effect of fine grained social distancing policies in real time and are very well suited to their purpose. We invite epidemiologists to provide feedback and collaborate to improve our model.

3   Our Model

3.1   Summary

To alleviate these problems we propose a different model structure. Our model inherits many features from agent-based simulation models but replaces the contacts between moving particles by contacts between individuals who work, go to school, live in a household and enjoy leisure activities. The structure of the model is depicted in Figure 1.
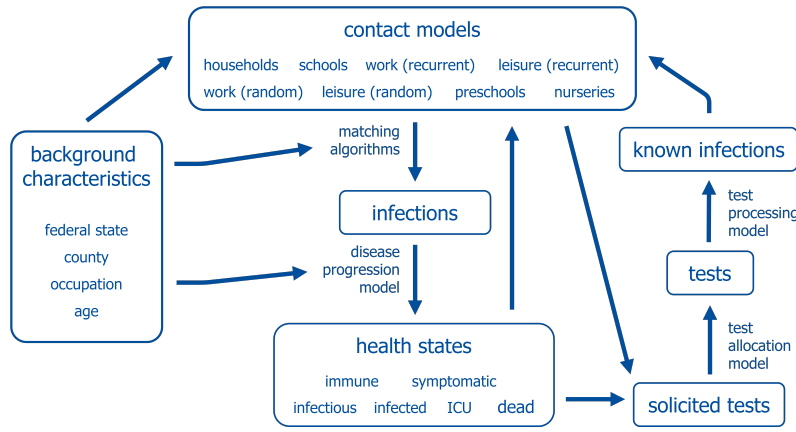
Figure 1: Simplified graph of the model

The background characteristics include age, county and occupation of each simulated individual. Contact models are functions that map individual characteristics into a predicted number of contacts. Currently we distinguish between eight types of contact models which are all listed in Figure 1: households, recurrent and random work contacts, recurrent and random leisure contacts, and nursery, preschool, and school contacts.

The predicted number of contacts is translated into infections by a matching algorithm. There are different matching algorithms for recurrent contacts (e.g. classmates, family members) and non-recurrent contacts (e.g. clients, contacts in supermarkets). The infection probability can differ for each contact type. All types of contacts can be assortative with respect to geographic and demographic characteristics.

Once a person is infected, the disease progresses in a fairly standard way which is also depicted in Figure 2. Asymptomatic cases and cases with mild symptoms are infectious for some time and recover eventually. Cases with severe symptoms additionally require hospitalization and lead to either recovery or death.

People who experience symptoms as well as people who had a risk contact can demand a test. Test allocation models that take into account the testing capacity determine which individuals actually receive a test. Test processing models determine in which order the tests are processed and when infected individuals are notified.

In addition, people who have symptoms, received a positive test, or had a risk contact can reduce their number of contacts across all contact types endogenously.

The model makes it very simple to translate policies into model quantities. For example, school closures imply the complete suspension of school contacts. A strict lockdown implies shutting down work contacts of all people who are not employed in a systemically relevant sector. It is also possible to have more sophisticated policies that condition the number of contacts on observable characteristics, risk contacts or health states.

Another key advantage of the model is that the number of contacts an individual has of each contact type can be calibrated from publicly available data (Mossong et al., 2008). This in turn allows us to estimate policy-invariant infection probabilities from time series of infection and death rates using the method of simulated moments (McFadden, 1989). Since the infection probabilities are time-

invariant, data collected since the beginning of the pandemic can be used for estimation. Moreover, since we can model the testing strategies that were in place at each point in time, we can correct the estimates for the fact that not all infections are observed.

In the following sections we describe each of the model components in more detail.

3.2 Modeling Numbers of Contacts

Consider a hypothetical population of 1,000 individuals in which 50 were infected with a novel infectious disease. From this data alone, it is impossible to say whether only those 50 people had contact with an infectious person and the disease has an infection probability of 1 in each contact or whether everyone met an infectious person but the disease has an infection probability of only 5 percent per contact. SEIR models do not even try to distinguish contact frequency from the infectiousness of each contact and combine the two in one parameter that is not invariant to social distancing policies.

To model social distancing policies, we need to disentangle the effects of the number of contacts of each individual and the effect of policy-invariant infection probabilities specific to each contact type. Since not all contacts are equally infectious, we distinguish different contact types.

The number and type of contacts in our model can be easily extended. Each type of contacts is described by a function that maps individual characteristics, health states and the date into a number of planned contacts for each individual. This allows to model a wide range of contact types.

Currently, there are the following contact types:

- Households: Each household member meets all other household members every day. The household sizes and structures are calibrated to be representative for Germany.

- Random non-work contacts: Each person has contacts with randomly drawn other people. This contact type reflects contacts during pure leisure activities as well as non leisure activities such as grocery shopping or medical appointments.

- Random work contacts: Each working adult has contact to randomly drawn other people.

- Recurrent daily work contacts: Each working adult meets other workers every day. This is meant to capture work colleagues.

- Recurrent weekly work contacts: Each working adult meets other workers once per week. We randomize over the days on which the meetings take place. This is meant to capture meetings with clients, superiors or other colleagues which happen infrequently.

- Schools: Each student meets all of his classmates every day. Class sizes are calibrated to be representative for Germany. Schools are closed on weekends and during vacations, which vary by states.

- Preschools: Similar to schools but for younger children.

- Nurseries: Similar to schools but for very young children.

The number of random and recurrent contacts at the workplace and at home is calibrated with data provided by Mossong et al., 2008. For details see Section 4. In particular, we sample the number of contacts or group sizes from empirical distributions that sometimes depend on age. Instead it would also be possible to use economic or other behavioral models to predict the number of contacts.

Theoretically, each contact type can have its own infection probability. However, to reduce the number of free parameters and thus avoid a potential over-fitting we impose some constraints. For now, infection probabilities in schools, preschools and nurseries are equal. Moreover, we restrict all work contacts to have the same infection probability.

3.3   Reducing Numbers of Contacts Through Policies

The main motivation of our model is to predict the effect of policies that affect the number of contacts people have. Examples range from school closures and lockdowns to more nuanced policies such as a mandatory quarantines for symptomatic individuals or a class splitting policy where only half of the students come to school in person and the other half joins digitally with weekly rotation.

Instead of thinking of policies as completely replacing how many contacts people have, it is often more helpful to think of them as adjusting the pre-pandemic number of contacts.

Therefore, we implement policies as a step that happens after the number of contacts is calculated but before individuals are matched.

On an abstract level, a policy is a functions that modifies the number of contacts of one contact type. For example, school closures simply set all school contacts to zero. A lockdown where only essential workers are allowed to work means that approximately two thirds of the working population have zero work contacts and the rest has the same number of contacts as before.

This, in conjunction with our fine grained contact types, allows us to easily implement a wide variety of policies. Allowing policies to depend on the health states of the entire population means that adaptive lockdowns where, for example, schools close when a certain threshold of infections is surpassed at the county level would be as simple as determining which counties are above the threshold and then setting all school contacts in these counties to zero.

The dependency of policies on health states also makes it possible to model contact tracing. For example, a policy could check whether each child has a classmate who's received a positive test result and then bar all children of that class from attending school.

Some policies can be easily implemented if the background characteristics are suitably extended. For example, a schooling policy with split classes, where each half attends school every other week can be implemented by storing whether the child would attend in even or odd weeks in the background characteristics and then using that information in the policy function.

For some policies the exact effect on each contact type is not easy to determine. If this refers to a policy during the estimation period, it is possible to estimate such parameters by fitting the model to time series data of infection rates. This is only possible if the policy was not active during the whole estimation period and thus the infection probabilities can be identified separately. If instead it refers to a policy that we want to simulate, we make a scenario analysis in which the model is simulated

with several assumptions about how the policy affects the number of contacts.

3.4   Endogenous Contact Reductions

Policies are not the only way in which the number of contacts are reduced compared to the pre-pandemic level. It is important to model those other channels. Otherwise, the effect of policies would be overestimated and policy recommendations based on the model would be biased.

Examples of endogenous contact reductions are manifold: symptomatic people stay at home; Members of risk groups try to reduce their number of contacts more strongly than others; People self-isolate if they know they had a risk contact.

Since we model the number of contacts as arbitrary functions of background characteristics and health states, it is easy to implement such considerations.

In our current empirical application we only model that symptomatic people reduce their number of contacts across all contact types (except for households) by 70 %. Within households they reduce contacts by 50%. We are working on extending this to allow for formal and informal contact tracing as well as quarantines after positive test results.

3.5   Matching Individuals

The empirical data described above only allows to estimate the number of contacts each person has. In order to simulate transmissions of Covid-19, the numbers of contacts has to be translated into actual meetings between people. This is achieved by matching algorithms:

As described in section 3.2, some contact types are recurrent (i.e. the same people meet regularly), others are non-recurrent (i.e. it would only be by accident that two people meet twice). The matching process is different for recurrent and non recurrent contact models.

Recurrent contacts are described by two components: 1) A variable in the background characteristics. An example would be a school class identifier which could come from actual data or be drawn randomly to achieve representative class sizes. 2) A deterministic or random function that takes the value 0 (non-participating) and 1 (participating) and can depend on the weekday, date and health state. This can be used to model vacations, weekends or symptomatic people who stay home (see section 3.4 for details).

The matching process for recurrent contacts is then extremely simple: On each simulated day, every person who does not stay home meets all other group members who do not stay home. The assumption that all group members have contacts with all other group members is not fully realistic, but seems like a good approximation to reality, especially in light of the suspected role of aerosol transmission for Covid-19 (Anderson et al., 2020; Morawska et al., 2020).

The matching in non-recurrent contact models is more difficult and implemented in a two stage sampling procedure to allow for assortative matching. Currently most contact models are assortative with respect to age (it is more likely to meet people from the same age group) and county (it is more likely to meet people from the same county) but in principle any set of discrete variables can be

used. This set of variables that influence matching probabilities introduce a discrete partition of the population into groups. The first stage of the two stage sampling process samples on the group level. The second stage on the individual level.

Below, we first show pseudo code for the non-recurrent matching algorithm and then describe how the algorithm works in words.

```
while unmatched contacts remain:
    model, id = draw contact model and individual
    for contact in remaining_contacts[id, model]:
        draw group of other person
        draw other person from that group
        determine if id infects other or vice versa
        remaining_contacts[id, model] -= 1
        remaining_contacts[other, model] -= 1
```

We first randomly draw a contact type and individual. For each contact oth the drawn contact type that person has, we first draw the group of the other person (first stage). Next, we calculate the probability to be drawn for each member of the group, based on the number of remaining contacts. I.e. people who have more remaining contacts are drawn with a higher probability. This has to be re-calculated each time because with each matched contact, the number of remaining contacts changes. We then draw the other individual, determine whether an infection takes place and if so update the health states of the newly infected person. Finally, we reduce the number of remaining contacts of the two matched individuals by one.

The recalculation of matching probabilities in the second stage is computationally intensive because it requires summing up all remaining contacts in that group. Using a two stage sampling process where the first stage probabilities remain constant over time makes the matching computationally much more tractable because the number of computations increases quadratically in the second stage group size.

3.6   Course of the Disease

The following medical parameters describing the progression of the disease are taken from systematic reviews (e.g. He et al., 2020). After an infection occurs, the disease progresses in the way depicted in Figure 2.

First, infected individuals will become infectious after one to five days. About one third of people stays asymptomatic. The rest develops symptoms about one to two days after they become infectious. Modeling asymptomatic and pre-symptomatic cases is important because those people do not reduce their contacts or demand a test and can potentially infect many other people (Donsimoni et al., 2020).

A small share of symptomatic people will develop strong symptoms that require intensive care. The exact share and time span is age-dependent. An age-dependent share of intensive care unit (ICU) patients will die after spending up to 32 days in intensive care. Moreover, if the ICU capacity was reached, all patients who require intensive care but do not receive it die.
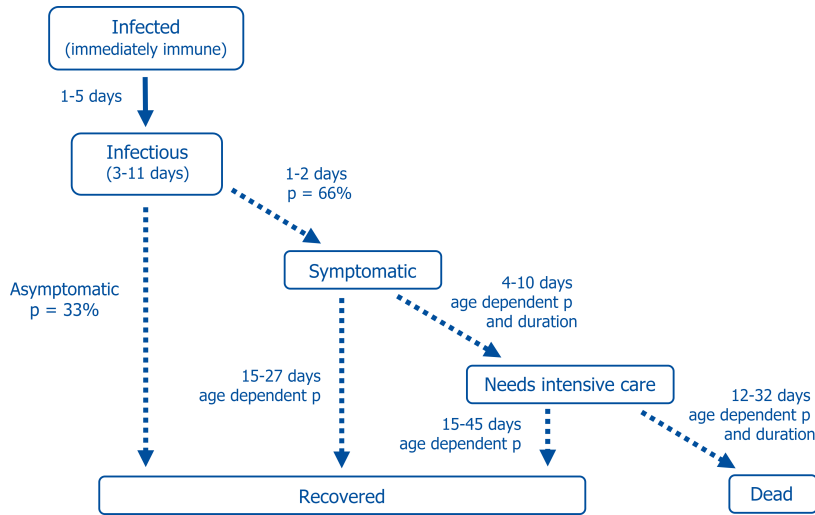
Figure 2: Course of Disease in the model

It would be easy to make the course of disease even more fine grained. For example, we could model people who require hospitalization but not intensive care. So far we opted against that because only the ICU capacity might become a bottleneck in Germany.

We allow the progression of the disease to be stochastic in two ways: Firstly, state changes only occur with a certain probability (e.g. only a fraction of infected individuals develops symptoms). Secondly, the number of periods for which an individual remains in a state is drawn randomly. The parameters that govern these processes are taken from the literature. They can vary with the age of an individual.

Detailed information on the calibration of the disease parameters is available as part of our online documentation.

4   Calibration and Estimation

The model is described by a large number of parameters that govern the number of contacts a person has, the likelihood of becoming infected on each contact, the likelihood of developing light or strong symptoms or even dying from the disease as well as the duration each stage of the disease takes.

Most of these parameters can be calibrated from existing datasets or the medical literature. Only the infection probabilities have to be estimated inside the model by fitting it to time series data of case numbers and fatality rates.

4.1   Medical Parameters

4.1.1   Length of Presymptomatic Stage / Incubation Period

Estimates of the incubation period usually give a range from 2 to 12 days. A meta analysis by McAloon et al., 2020 comes to the conclusion that "The incubation period distribution may be modelled with a lognormal distribution with pooled μ and σ parameters (95 % CIs) of 1.63 (95 % CI 1.51 to 1.75) and 0.50 (95 % CI 0.46 to 0.55), respectively." For simplicity we discretize this

distribution into four bins.

4.1.2 Begin of Infectiousness

The period between infection and onset of infectiousness is called latent or latency period. However, the latency period is rarely given in epidemiological reports on Covid-19. Instead, scientists and agencies usually report the incubation period, the period from infection to the onset of symptoms. A few studies used measurements of virus shedding to estimate infectiousness during the course of the disease. When measurements started before the onset of symptoms the development of the viral load before symptoms gives us an indication of number of days between the onset of infectiousness and symptoms.

The European Centre for Disease Prevention and Control estimates that people become infectious between one and two days before the symptoms set in. This is similar to He et al., 2020 who estimate this to take 2.3 days and is in line with Peak et al., 2020.

Given these numbers and the length of the incubation period we can calculate the latency period for symptomatic people. To our knowledge no estimates for the latency period of asymptomatic cases of COVID-19 exist. We assume it to be the same for symptomatic and asymptomatic cases.

Thus, we arrive at the following distribution for latency periods: 40% have one day. 35% have two days. 20% have three days and 5% have 5 days.

4.2 Duration of Infectiousness

We assume that the duration of infectiousness is the same for both symptomatic and asymptomatic individuals as evidence suggests little differences in the transmission rates of SARS-CoV-2 virus between symptomatic and asymptomatic patients (Yin and Jin, 2020) and that the viral load between symptomatic and asymptomatic individuals are similar (Zou et al., 2020, Byrne et al., 2020, Singanayagam et al., 2020).

Our distribution of the duration of infectiousness is based on Byrne et al., 2020.

For symptomatic cases they arrive at 0-5 days before symptom onset (figure 2) and 3-8 days of infectiousness afterwards.[1] Thus, we arrive at 0 to 13 days as the range for infectiousness among individuals who become symptomatic (see also figure 5). This duration range is very much in line with the meta-analysis' reported evidence for asymptomatic individuals (see their figure 1). Thus, we arrive at 0 to 13 days as the range for infectiousness among individuals who become symptomatic. This duration range is very much in line with the meta-analysis' reported evidence for asymptomatic individuals.

Following this evidence we assume the following discretized distribution of the infectiousness period: 10% of individuals are infectious for three days, 25% for five days, another 25% for seven days, 20% for nine days and 20% for eleven days.

---

[1]Viral loads may be detected much later but 8 days seems to be the time after which most people are culture negative, as also reported by Singanayagam et al., 2020

4.3 Duration of Symptoms

We use the duration to recovery of mild and moderate cases reported by Bi et al., 2020, Figure S3, Panel 2 for the duration of symptoms for asymptomatic and non-ICU requiring symptomatic cases.

We collapse the data to the following distribution: 10% recover after 15 days and 30% require 18, 22 or 27 days respectively.

These numbers are only used for mild cases. We do not disaggregate by age. Note that the length of symptoms is not very important in our model given that individuals stop being infectious before their symptoms cease.

4.3.1 Time from Symptom Onset to Admission to ICU

The data on how many percent of symptomatic patients will require ICU is pretty thin. We rely on data by the US CDC (Stokes et al., 2020) and the OpenABM-Project (2020-09-14).

From these sources we arrive at the following probabilities of requiring intensive care:

Table 1: Shares of symptomatic patients who will require ICU care by age groups.

Age Group	Share
0-9	0.00005
10-19	0.00030
20-29	0.00075
30-39	0.00345
40-49	0.01380
50-59	0.03404
60-69	0.10138
70-79	0.16891
80-100	0.26871

For those who will require intensive care we follow Chen et al., 2020 who estimate the time from symptom onset to ICU admission as 8.5 ± 4 days.

This aligns well with numbers reported for the time from first symptoms to hospitalization: Gaythorpe et al., 2020 report a mean of 5.76 with a standard deviation of 4. This is also in line with the durations collected by the Robert Koch Institut.

We assume that the time between symptom onset and ICU takes 4, 6, 8 or 10 days with equal probabilities. These times mostly matter for the ICU capacities.

4.3.2 Death and Recovery from ICU

We take the survival probabilities and time to death and time until recovery from intensive care from the OpenABM Project.

They report time until death to have a mean of 11.74 days and a standard deviation of 8.79 days. Approximating this with the normal distribution, we have nearly 10% probability mass below 0. We

use it nevertheless as several other distributions (such as chi squared and uniform) were unable to match the variance. Discretizing this leads to 41% of individuals who die from Covid-19 to die after one day in intensive care. 22% day after 12 days, 29% after 20 days and 7% after 32 days. Again, we rescale this for every age group among those that will not survive.

They report time until recovery to have a mean of 18.8 days and a standard deviation of 12.21 days. Approximating this with the normal distribution, we have over 5% probability mass below 0. Discretizing this of those who recover in intensive care 22% do so after one day, 30% after 15 days, 28% after 25 days and 18% after 45 days.

4.4 Number of Contacts

We calibrate the parameters for the predicted numbers of contacts from contact diaries of over 2000 individuals from Germany, Belgium, the Netherlands and Luxembourg (Mossong et al., 2008). Each contact diary contains all contacts an individual had throughout one day, including information on the other person (such as age and gender) and information on the contact. Importantly, for each contact individuals entered of which type the contact (school, leisure, work etc.) was and how frequent the contact with the other person is.

Thus, we can use the empirical distributions from this data as pre-pandemic number of contacts.

4.5 Assortative Matching

As mentioned in section 3.5, the probability that two individuals are matched can depend on background characteristics. In particular, we allow this probability to depend on age and county of residence. While we do not have good data on geographical assortativeness and just roughly calibrate it such that 60 % of contacts are within the same county,[2] we can calibrate the age assortativeness from the same data we use to calibrate the number of contacts.

|	0-9	10-19	20-29	30-39	40-49	50-59	60-69	70-79	80-100
0-9 (697 obs.)	0.3	0.1	0.1	0.1	0.1	0.1	0.1	0.0	0.0
10-19 (1309 obs.)	0.1	0.6	0.1	0.0	0.1	0.0	0.0	0.0	0.0
20-29 (609 obs.)	0.0	0.1	0.5	0.1	0.1	0.1	0.0	0.0	0.0
30-39 (401 obs.)	0.1	0.1	0.2	0.3	0.2	0.1	0.1	0.0	0.0
40-49 (433 obs.)	0.0	0.1	0.1	0.2	0.3	0.1	0.1	0.1	0.0
50-59 (452 obs.)	0.1	0.0	0.1	0.1	0.2	0.2	0.1	0.1	0.1
60-69 (410 obs.)	0.0	0.0	0.1	0.1	0.2	0.2	0.2	0.1	0.1
70-79 (141 obs.)	0.0	0.0	0.0	0.2	0.1	0.1	0.2	0.2	0.0
80-100 (31 obs.)	0.0	0.0	0.0	0.0	0.0	0.1	0.2	0.5	0.0

age_group_of_contact

Figure 3: Distribution of random non-work contacts by age of participants

Figure 3 shows that overall contacts within age groups or age groups that are closer to each other

---

[2]We are working on improving this estimation with mobility data.

are more frequent.

4.6 Infection Probabilities

To calibrate infection probabilities outside of the model, it would be important to know the exact duration and distance of each contact type as well as virus loads. Since this is not available in any dataset, we instead estimate those parameters inside the model by minimizing the distance between simulated and observed infection rates. Since our model includes a lot of randomness, we average simulated infection rates over several model runs.

Currently, we use data for Germany from August to October (both inclusive). We do not use earlier periods to save computational time. Moreover, we would be worried that the infection probabilities have small seasonal variation that we currently cannot model. However, we plan to expand the estimation period soon.

To avoid overfitting and simplify the numerical optimization problem, we only allow for four different probabilities: 1) for contacts in schools, preschools and nurseries. 2) for work contacts. 3) for households. 4) for leisure activities.

5 Model Validation

We validate our model in two ways: 1) We look at the in-sample fit over the estimation period. 2) We look at the out-of-sample fit for November. The last one is a challenging test for our model because there was a strong policy change between the estimation period and November. The model convincingly passes both tests.

5.1 In-Sample Fit

Despite fitting only four free parameters, the in-sample fit is very good. The best fit is achieved in the largest age groups. This is so mechanically, because we weight the deviations between simulated and observed infection rates by group sizes. The worst fit is achieved for the 80 to 100 years old. The model predicts too few infections for these groups because they have very few contacts in all contact types we have included so far. We plan to address this issue soon by adding another contact type that captures all contacts in the data by Mossong et al., 2008 that we have not included so far. Moreover, we expect an improved model fit when we allow for more free parameters.

5.2 Out-of-Sample Fit

We can assess the out-of-sample fit by projecting the effect of the lockdown light and comparing it to case numbers until now. It is important to note that this is not just a simple extrapolation of a time trend because the lockdown light only started after the estimation period. The out-of-sample fit can be assessed in Figure 5.

The model correctly predicts the effect of the lockdown light with reasonable accuracy. In particular,
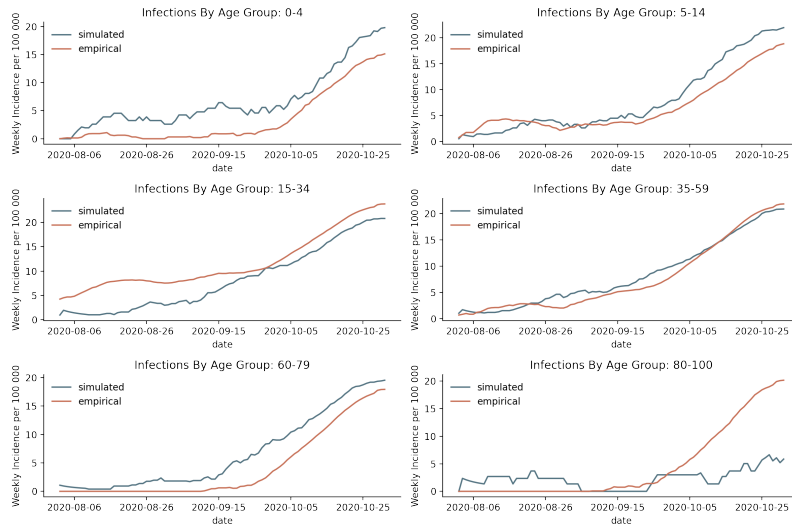
Figure 4: Actual vs. simulated infection and fatality rates


Figure 5: Projected effect of the lockdown-light

the actual case numbers are between our neutral and pessimistic projection. The plot also shows that ending the lockdown light as planned on November 30 would lead to an explosive growth in case numbers in all scenarios.

6 Conclusion

We propose a simulation based model of infectious disease transmission that is designed to predict the effects of fine grained social distancing policies. In particular, the model can be used to model policies such as several ways of splitting school classes or work reduction policies that affect essential and non-essential workers differently. Both policies would be hard to implement in standard SEIR or agent based simulation models.

To predict the effects of such policies, it is not only important to have a way of expressing such flexible policies in terms of model quantities, but also to incorporate heterogeneity in disease progression as well as meeting patterns. We calibrate age dependent disease progression parameters from the medical literature and age dependent contact frequencies from contact diaries. Moreover, we distinguish eight types of contacts out of which some are only relevant for certain age groups.

The model has a good fit on past German case numbers and passes an out of sample validation despite a drastic change in the policy environment between the estimation period and the validation period.

Despite these encouraging results we still see the model as work in progress and plan to implement more features such as a detailed model of testing and contact tracing. Moreover, the estimation of the infection probabilities and the model fit will improve as more data becomes available.

We invite researchers from any discipline, but particularly epidemiologists to provide feedback on the model and welcome collaborations.

References

Acemoglu, Daron, Victor Chernozhukov, Iván Werning, and Michael D Whinston (May 2020). *Optimal Targeted Lockdowns in a Multi-Group SIR Model*. Working Paper 27102. National Bureau of Economic Research. DOI: 10.3386/w27102. URL: http://www.nber.org/papers/w27102.

Anderson, Elizabeth L., Paul Turnham, John R. Griffin, and Chester C. Clarke (May 2020). "Consideration of the Aerosol Transmission for COVID-19 and Public Health". In: *Risk Analysis* 40.5, pp. 902–907. DOI: 10.1111/risa.13500.

Avery, Christopher, William Bossert, Adam Clark, Glenn Ellison, and Sara Fisher Ellison (Nov. 2020). "An Economist's Guide to Epidemiology Models of Infectious Disease". In: *Journal of Economic Perspectives* 34.4, pp. 79–104. DOI: 10.1257/jep.34.4.79. URL: https://www.aeaweb.org/articles?id=10.1257/jep.34.4.79.

Berger, David W, Kyle F Herkenhoff, and Simon Mongey (Mar. 2020). *An SEIR Infectious Disease Model with Testing and Conditional Quarantine*. Working Paper 26901. National Bureau of Economic Research. DOI: 10.3386/w26901. URL: http://www.nber.org/papers/w26901.

Bi, Qifang et al. (Mar. 2020). "Epidemiology and Transmission of COVID-19 in Shenzhen China: Analysis of 391 cases and 1,286 of their close contacts". In: DOI: 10.1101/2020.03.03.20028423.

Byrne, Andrew William et al. (Aug. 2020). "Inferred duration of infectious period of SARS-CoV-2: rapid scoping review and analysis of available evidence for asymptomatic and symptomatic COVID-19 cases". In: *BMJ Open* 10.8, e039856. DOI: 10.1136/bmjopen-2020-039856.

Chen, Jun et al. (May 2020). "Clinical progression of patients with COVID-19 in Shanghai, China". In: *Journal of Infection* 80.5, e1–e6. DOI: 10.1016/j.jinf.2020.03.004.

Cuevas, Erik (June 2020). "An agent-based model to evaluate the COVID-19 transmission risks in facilities". In: *Computers in Biology and Medicine* 121, p. 103827. DOI: 10.1016/j.compbiomed.2020.103827. URL: https://doi.org/10.1016/j.compbiomed.2020.103827.

Donsimoni, Jean Roch, René Glawion, Bodo Plachter, and Klaus Wälde (2020). "Projecting the spread of COVID-19 for Germany". In: *German Economic Review* 21.2, pp. 181–216. DOI: https://doi.org/10.1515/ger-2020-0031. URL: https://www.degruyter.com/view/journals/ger/21/2/article-p181.xml.

Dorn, Florian, Janos Gabler, Hans-Martin von Gaudecker, Andreas Peichl, Tobias Raabe, and Klara Rörhl (2020a). "Wenn Menschen (keine) Menschen treffen: Simulation der Auswirkungen von Politikmaßnahmen zur Eindämmung der zweiten Covid-19-Welle". In: *ifo Schnelldienst Digital*. URL: https://www.ifo.de/publikationen/ifo-schnelldienst.

Dorn, Florian, Sahamoddin Khailaie, Marc Stöckli, Sebastian Binder, Berit Lange, Patrizio Vanella, Timo Wollmershäuser, Andreas Peichl, Clemens Fuest, and Michael Meyer-Hermann (2020b). "Das gemeinsame Interesse von Gesundheit und Wirtschaft: Eine Szenarienrechnung zur Eindämmung der Corona- Pandemie". ger. In: *ifo Schnelldienst Digital* 1.6. ISSN: 2700-8371. URL: http://hdl.handle.net/10419/223322.

Gaythorpe, K et al. (2020). "Report 8: Symptom progression of COVID-19". In: DOI: 10.25561/77344.

Grimm, Veronika, Friederike Mengel, and Martin Schmidt (2020). "Extensions of the SEIR Model for the Analysis of Tailored Social Distancing and Tracing Approaches to Cope with COVID-19". In: *medRxiv*. DOI: 10.1101/2020.04.24.20078113. eprint: https://www.medrxiv.org/content/early/2020/04/29/2020.04.24.20078113.full.pdf. URL: https://www.medrxiv.org/content/early/2020/04/29/2020.04.24.20078113.

He, Xi, Eric HY Lau, Peng Wu, Xilong Deng, Jian Wang, Xinxin Hao, Yiu Chung Lau, Jessica Y Wong, Yujuan Guan, Xinghua Tan, et al. (2020). "Temporal dynamics in viral shedding and transmissibility of COVID-19". In: *Nature medicine* 26.5, pp. 672–675.

Hinch, Robert et al. (Sept. 2020). "OpenABM-Covid19 - an agent-based model for non-pharmaceutical interventions against COVID-19 including contact tracing". In: DOI: 10.1101/2020.09.16.20195925.

McAloon, Conor et al. (Aug. 2020). "Incubation period of COVID-19: a rapid systematic review and meta-analysis of observational research". In: *BMJ Open* 10.8, e039652. DOI: 10.1136/bmjopen-2020-039652. URL: https://doi.org/10.1136/bmjopen-2020-039652.

McFadden, Daniel (1989). "A method of simulated moments for estimation of discrete response models without numerical integration". In: *Econometrica: Journal of the Econometric Society*, pp. 995–1026.

Morawska, Lidia et al. (Sept. 2020). "How can airborne transmission of COVID-19 indoors be minimised?" In: *Environment International* 142, p. 105832. DOI: 10.1016/j.envint.2020.105832.

Mossong, Joël, Niel Hens, Mark Jit, Philippe Beutels, Kari Auranen, Rafael Mikolajczyk, Marco Massari, Stefania Salmaso, Gianpaolo Scalia Tomba, Jacco Wallinga, et al. (2008). "Social contacts and mixing patterns relevant to the spread of infectious diseases". In: *PLoS medicine* 5.3.

Peak, Corey M, Rebecca Kahn, Yonatan H Grad, Lauren M Childs, Ruoran Li, Marc Lipsitch, and Caroline O Buckee (Sept. 2020). "Individual quarantine versus active monitoring of contacts for the mitigation of COVID-19: a modelling study". In: *The Lancet Infectious Diseases* 20.9, pp. 1025–1033. DOI: 10.1016/s1473-3099(20)30361-3. URL: https://doi.org/10.1016/s1473-3099(20)30361-3.

Silva, Petrônio C.L., Paulo V.C. Batista, Hélder S. Lima, Marcos A. Alves, Frederico G. Guimarães, and Rodrigo C.P. Silva (Oct. 2020). "COVID-ABS: An agent-based model of COVID-19 epidemic to simulate health and economic effects of social distancing interventions". In: *Chaos, Solitons & Fractals* 139, p. 110088. DOI: 10.1016/j.chaos.2020.110088. URL: https://doi.org/10.1016/j.chaos.2020.110088.

Singanayagam, Anika, Monika Patel, Andre Charlett, Jamie Lopez Bernal, Vanessa Saliba, Joanna Ellis, Shamez Ladhani, Maria Zambon, and Robin Gopal (Aug. 2020). "Duration of infectiousness and correlation with RT-PCR cycle threshold values in cases of COVID-19, England, January to May 2020". In: *Eurosurveillance* 25.32. DOI: 10.2807/1560-7917.es.2020.25.32.2001483.

Stokes, Erin K., Laura D. Zambrano, Kayla N. Anderson, Ellyn P. Marder, Kala M. Raz, Suad El Burai Felix, Yunfeng Tie, and Kathleen E. Fullerton (June 2020). "Coronavirus Disease 2019 Case Surveillance — United States, January 22–May 30, 2020". In: *MMWR. Morbidity and Mortality Weekly Report* 69.24, pp. 759–765. DOI: 10.15585/mmwr.mm6924e2.

Yin, Guosheng and Huaqing Jin (May 2020). "Comparison of Transmissibility of Coronavirus Between Symptomatic and Asymptomatic Patients: Reanalysis of the Ningbo COVID-19 Data". In: *JMIR Public Health and Surveillance* 6.2, e19464. DOI: 10.2196/19464.

Zou, Lirong et al. (Mar. 2020). "SARS-CoV-2 Viral Load in Upper Respiratory Specimens of Infected Patients". In: *New England Journal of Medicine* 382.12, pp. 1177–1179. DOI: 10.1056/nejmc2001737.