MILLIMAN WHITE PAPER

Pandemic risk modelling in Solvency II internal models: Example of COVID-19

Alexandre Boumezoued
Eve Titon

Milliman

The current pandemic due to COVID-19 raises many questions regarding the modelling of the propagation of a new virus, with specific characteristics, and with incomplete historical data. In this context of epidemics in an early stage, estimating the key parameters of the dynamics (infection rate, fatality rate) remains challenging.

In this paper we discuss actuarial perspectives for the calculation of the capital requirement related to catastrophic pandemic mortality as required by the Solvency II regulation. We also provide an overview of possible models to consider and explore some biases and challenges in parameter estimation.

Pandemic risk: Beyond frequency-severity models

Under the Solvency II regulation, the life catastrophe risk stems from extreme death events that are not sufficiently captured by the mortality risk sub-module. These are one-time shocks from the extreme, adverse tail of the probability distribution that are not adequately represented by extrapolation from more common events and for which it is usually difficult to specify a loss value, and thus an amount of capital to hold, see [1].

The regulation especially focusses on examples such as a contagious disease process or a pandemic which can affect many people simultaneously. It is recalled that the Solvency II catastrophe mortality stress is an absolute increase (additive value on top of base mortality rates) of 0.15%, which is considered as a 1 in 200 year pandemic stress.

In France, for example, approximately 600,000 deaths a year are observed, while the Solvency II pandemic shock would result in approximately 100,000 excess deaths. This number could be seen as high, but is in fact below current pessimistic scenarios of COVID-19 if no mitigation measures are taken.

LIMITS OF FREQUENCY/SEVERITY MODELS

The calibration of the Solvency II stress was originally based on an epidemiological model. However, the final representation of the stress relies on consideration of frequency (1:200) and severity (0.15%). Moreover, the absolute increase in mortality rates does not depend on the age considered, at first order, as well as, at second order, the health condition, whereas the pandemic analyses for COVID-19 do show a clear differential with age and comorbidity.

The development of Solvency II internal models in Europe covered that of the pandemic risk module from early on. A significant part of those models was initially (or is still) based on frequency/severity models. One classical approach is to estimate the components separately:

- The frequency component, looking at past pandemic occurrences; as a core example, the original Spanish flu has been used to estimate a return period with an order of magnitude between 1:100 and 1:200 in some studies.
- As for the severity component, a possible approach is to try to determine what the impact would be of such past pandemics in current conditions. Those can be made of two opposite effects: first, due to medical improvement, the mortality rate would reduce at each age (noticing that overall impact can be high in an ageing population); however, increased transportation flows would increase contamination and the spread of the pandemics.

The current experience on COVID-19 reminds us that more structural models are needed to appropriately capture the risk of propagation and the impact of pandemics. Some (re)insurance companies have taken this direction and it is believed that there will be greater movement to implement more complex models, encouraged by the regulation.

WHAT A GOOD INTERNAL MODEL FOR PANDEMIC RISK WOULD BE

In the light of the current COVID-19 pandemic, one can describe some key features which could be expected in the near future by the regulator of a “good” internal model for pandemic mortality risk. We provide a list of such features below:

- Model the spread of the epidemics based on the contamination rate, adjusted by the incubation and contagion periods; population density in worldwide areas should be taken into account.
- Differentiate fatality rates per age and health/comorbidity conditions, and increase those rates in the event that the number of infected people exceeds hospitals’ capacity.

- Take into account mitigation measures and their timing, such as limitation in transportation flows, quarantine, development of emergency treatments, large scale testing approaches for early anticipated care.
- Transpose the impacts of the pandemic at national level to the insured population, by taking into account its specific age and characteristics features (medical selection, location to adjust for local healing capacities) and obviously the sums at risk.

As such, developing an appropriate pandemic risk model remains challenging because many components are interacting dynamically in the forecast and depend on a large set of assumptions and parameters, which can moreover evolve in the early stage of the pandemic. In the following, we provide an overview of possible models to consider and illustrate some biases and challenges in parameter estimation.

Note that the usefulness of a suitable model goes beyond the strict framework of regulatory capital calculation. It can also be used for internal risk studies and helps as a supporting decision tool for reinsurance valuation.

Modelling virus propagation

SIR-type models conventionally model epidemics. These are compartmental models in which each type of population, usually denoted by S-Susceptible, I-Infectious, R-Removed, interact with each other.

One of the simplest SIR-model is presented in Figure 1: because populations S and I are assumed to be in contact, each individual from population I infects individuals of population S with a so-called infection or contamination rate denoted by α. This means that at each time step the number of infected individuals is increased by αSI. As such, the model is said to allow for interactions, or to be non-linear.

Also, infected individuals are "removed"-- that is, recover or die, with a rate denoted by β and usually (ambiguously) named "recovery" rate, so that it decreases the number of infected individuals by βI at each time step. Note again that in this simple model, the population R contains both recoveries and deaths. Therefore, in this model, the number of infected lives evolves over time as a combination of the contamination rate and the recovery rate; the number of infected lives decreases with the recovery rate because a removed individual cannot infect others, but increases with the contamination rate itself.

FIGURE 1: SIR MODEL

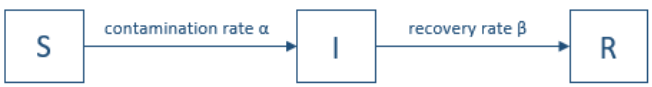

In the light of advanced characteristics of pandemic spread, such a simple model cannot fully replicate all the specificities of COVID-19. In Figure 2, a more sophisticated model is presented, with the following features:

- Infectious people are divided into two categories. I_{AS} are asymptomatic infectious people, and I are infectious people who present symptoms of the virus. Depending on the pandemics, asymptomatic infected people can or cannot infect other individuals. Moreover, precise data concerning the proportion of asymptomatic infected people among the population is often challenging to get, and assumptions are usually made.
- Recoveries are divided into R-recovered and D-deaths. Therefore the β parameter used in Figure 1 is to be split into a recovery rate and a death/fatality rate. Whereas the death rate can be found in the literature, the recovery rate (as the number of healthy recoveries among infected, per time unit) is usually available to a lesser extent given the precise information needed on the time-to-recovery duration.
- Some research suggested that individuals who recover from COVID-19 may be infected again, which is represented by the grey dotted arrow in Figure 2. Although no quantitative information about this phenomenon seems available for the moment, such a model can be useful to test possible under-estimation of the pandemic spread if one omits this possibility.

FIGURE 2: AN EXAMPLE OF EXTENDED SIR MODEL

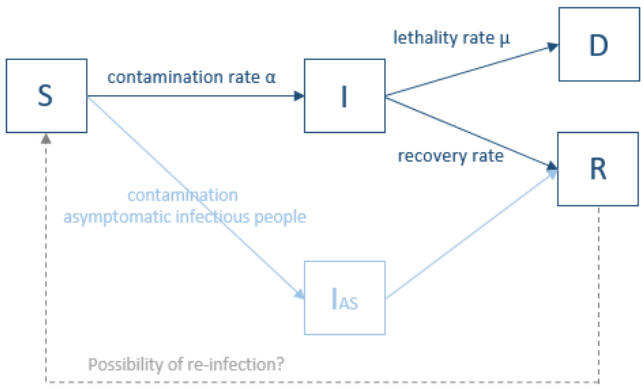

Under this more complex model, additional parameters remain to be estimated:

- In priority, the death rate, which is of particular interest in the remaining of this paper.
- The contamination rate leading to asymptomatic individuals, as well as the rate at which individuals who have recovered can be put into the S population to be potentially infected again. By nature, these quantities are challenging to estimate and assumptions have to be made.

To this model, one could add an incubation period, modelled by an additional compartment between S and I. Travels also could be represented by distinguishing infectious people depending on their propensity to travel or not. Mitigation actions such as quarantine or vaccination can be taken into account by making adjustments over time to the propagation rate.

CONTAMINATION RATE

The contamination rate α is difficult to estimate. Epidemiologists are rather focusing on the so-called basic reproduction number denoted by R_0, which is defined as the expected number of secondary cases produced by a single infection in a completely susceptible population. In the classical SIR model where population is scaled to one unit, then $R_0 = \frac{\alpha}{\beta}$.

That is, the reproduction number is the ratio between the contamination rate and the recovery rate (healthy recovery or death).

If $R_0 S(0) \leq 1$, where $S(0)$ is the initial size of the susceptible population, then an infected individual contaminates on average less than one susceptible individual, and therefore the disease does not spread. On the contrary, if $R_0 S(0) > 1$, then an infected individual contaminates on average more than one susceptible individual, and therefore the number of infected increases. Thus, the pandemic appears. Several publications about the basic reproduction number for COVID-19 provide estimates between 2 and 3.

The modelling should allow R_0 to vary over time (mitigation measures to stem the epidemic can be taken, for instance, to reduce α), and among the different countries. For example, reference [2] gives an estimation over time of the basic reproduction number for more than 20 countries.

FATALITY RATE

The case fatality rate μ is at the heart of all the debates today: is the calculation used and communicated under- or over-estimated? The next section is dedicated to the discussion of this point.

Calculation of the fatality rate

CLASSICAL ESTIMATOR

The death rate is commonly called case fatality rate (CFR). It represents the risk of death among infected cases.

Let us denote by C_t the cumulative number of confirmed cases and D_t the cumulative numbers of deaths on day t of the epidemic: a simple formula for the CFR is $\mu_t = \frac{D_t}{C_t}$.

SEVERAL BIASES

However, this calculation may present several important biases:

- The asymptomatic infectious cases are not detected as well as other symptomatic cases (because of a shortage of tests, for instance). This leads to an overestimation of the death rate because the exposure to risk is underestimated.
- The second bias, which tends to underestimate the risk, is of particular importance in the early stage of an epidemic. Among observed infected cases, some deaths that will occur have not yet been observed. This is a classical bias in survival analysis known as right censoring. That is why the delay in time from illness onset to death must be taken into account in the exposure calculation-- to give less importance to cases which have been very recently infected.

FOCUS ON UNDETECTED CASES

As mentioned before, the exposure to risk may be underestimated:

- In some countries, deaths are systematically tested for COVID-19. If the test is positive but the individual had not been detected before, the death is counted in D_t, but this individual was not counted in the exposure C_s for past calculations at time $s \leq t$.
- Asymptomatic infected people are not detected, and some symptomatic infected individuals are undetected, too.

Therefore, it is not a real death rate because the exposure does not take into account all infected people. Rather it is the probability of death of more severely infected individuals. It may not be an issue in the SIR modelling, if α is the "severe" contamination rate-- that is to say the transition rate from "Susceptible" to "severe symptomatic Infected". This could allow correction of the bias of underestimation of the exposure from the case fatality rate.

ADJUSTED CASE FATALITY RATE

The issue with bias due to deaths that have not yet been observed is commonly a right censoring issue, which arises in the early stage of an epidemic. At this point we need to introduce the probability density function of the time from illness onset to death (the relative frequency of the time from illness onset to death among fatal cases), denoted by f_s.

Thus, to recover an unbiased estimate, see [3], the cumulative number of infected cases $C_t = \sum_{s=1}^{t-1} c_{t-s}$ in the denominator is replaced by $\sum_{s=1}^{t-1} c_{t-s} f_s$, where c_t is the daily number of cases:

$$\mu'_t = \frac{D_t}{\sum_{s=1}^{t-1} c_{t-s} f_s} \geq \frac{D_t}{\sum_{s=1}^{t-1} c_{t-s}} = \mu_t.$$

This formula allows us to take into account the fact that a person who has been infected for only two days is less likely to die than someone who has been infected for 12 days, for example. Therefore, the first person contribution to the exposure to risk is less than that of the second one.

However, it is not simple to calibrate f because of the lack of public and precise data, so one can use a simplified estimate, proposed in [4], for example. As the time between illness onset and death ranges from 14 days to 8 weeks, the death rate can be estimated by dividing the number of deaths by the number of infected cases with illness onset older than 14 days. That is to say: for s from 1 to 14, $f_s = 0$, and for $s > 14$, $f_s = 1$. Such a simple rule can be applied given the data at hand for the daily number of cases, and leads to increased current estimates of mortality within symptomatic detected patients.

From national to insured population

Pandemics such as COVID-19 have an impact on life insurance because of the additional mortality caused by the virus. If the disease-related mortality rate is calculated at the national level, it must be adapted to the level of an insurance portfolio.

The insured population usually corresponds to the highest-income people. In order to target this population, the level of education is generally used as a proxy for income: higher education is correlated with higher socioeconomic class. One usually observes a reduction of global mortality in this sub-population, which represents the insured population. Note that as for pandemic, the mortality could be considered as similar for each age, but can also be thought to vary if one sees a link between socioeconomic status and health condition.

The transition from general to insured population has indeed to take into account pre-existing medical conditions. For COVID-19, it is well-known that the case fatality risk is higher in the parts of the population already suffering from cardiovascular diseases, diabetes, chronic respiratory diseases, hypertension and cancer. Thus, the use of statistics from medical underwriting will be useful to calculate the proportion of insured individuals suffering from these diseases. Also, when good health has been established at underwriting, the assumption of absence of comorbidity can be made to last one to three years after underwriting, for example.

Finally, the sums at risk are, of course, of particular importance, first because they provide information on previously cited covariates as a proxy (socioeconomic category) and also because portfolio heterogeneity will allow you to calculate confidence intervals for the expected outcome.

Modelling the unexpected

It is easier to describe an appropriate modelling framework based on known virus characteristics. It is more difficult to build a model able to capture unexpected features of viruses which could arise in the future. Some are listed below as possible ideas for modelling improvements:

- A virus could be tied to geographic-specific conditions that are favorable to its spread; this can include temperature, insect/animal species, food habits, human genes, etc. As such, special attention should be given to the inclusion of detailed geographic-specific features in the model.
- The reasons for propagation of the virus could be less clear to model, as for example contamination by food or insects; in the latter case, for example, population density of insects should be taken into account as a new variable to explain infection probability.

Finally, one can draw specific attention to the so-called “second order deaths” which are collateral to the pandemic crisis. For example, this can include people with unrelated serious conditions who need treatment but are deprioritised at some point. If the crisis reaches a severe threshold, then these effects become non-negligible and should be taken into account accordingly.

References

- [1] EIOPA-14-322, 25 July 2014. The underlying assumptions in the standard formula for the Solvency Capital Requirement calculation.
- [2] Sam Abbott, Joel Hellewell, James D Munday, June Young Chun, Robin N. Thompson, Nikos I Bosse, Yung-Wai Desmond Chan, Timothy W Russell, Christopher I Jarvis, CMMID nCov working group, Stefan Flasche, Adam J Kucharski, Rosalind Eggo & Sebastian Funk. 2020. *Temporal variation in transmission during the COVID-19 outbreak.* CMMID.
- [3] Tetsuro Kobayashi, Sung-mok Jung , Natalie M. Linton, Ryo Kinoshita, Katsuma Hayashi, Takeshi Miyama, Asami Anzai, Yichi Yang, Baoyin Yuan, Andrei R. Akhmetzhanov, Ayako Suzuki & Hiroshi Nishiura. 2020. Communicating the Risk of Death from Novel Coronavirus Disease (COVID-19). *Journal of clinical medicine.*
- [4] David Baud, Xiaolong Qi, Karin Nielsen-Saines, Didier Musso, Léo Pomar & Guillaume Favre. 2020. *Real estimates of mortality following COVID-19 infection.* Available at https://www.thelancet.com/action/showPdf?pii=S1473-3099%2820%2930195-X

In a crisis, clarity can be elusive. When the new normal bears little resemblance to the past, and data and information for decision making are in short supply, organizations turn to their trusted advisors to help them find a way forward. For more than 70 years, Milliman has been among the most trusted of those advisors to the Life insurance industry.

Our ongoing efforts are helping our clients understand, anticipate, and respond to the full range of possible impacts from this public health crisis. Milliman is advising the full spectrum of stakeholders to help them answer important business questions. We are:

- helping insurers make certain they have **adequate financial reserves and sufficient capital**
- supporting insurers with the adjustment of their **enterprise risk management and ORSA frameworks**
- measuring how this crisis affects **life insurance products' design, pricing, valuation, and reinsurance**

The need for professional insight and advice is now more important than ever. For more information about how Milliman can help your organization find clarity in this time of uncertainty, contact your Milliman consultant or email us at COVID19@milliman.com.

Milliman

Milliman is among the world's largest providers of actuarial and related products and services. The firm has consulting practices in life insurance and financial services, property & casualty insurance, healthcare, and employee benefits. Founded in 1947, Milliman is an independent firm with offices in major cities around the globe.

milliman.com

CONTACT

Alexandre Boumezoued
alexandre.boumezoued@milliman.com

Eve Titon
eveelisabeth.titon@milliman.com