

The Influenza A(H1N1)v Pandemic: An Exploratory System Dynamics Approach

Erik Pruyt and Caner Hamarat

Faculty of Technology, Policy & Management; Delft University of Technology
P.O. Box 5015, 2600 GA Delft, The Netherlands – E-mail: e.pruyt@tudelft.nl

10 June 2010*

Abstract

This paper presents a small exploratory System Dynamics model related to the dynamics of the 2009 flu pandemic, also known as the Mexican flu, swine flu, or A(H1N1)v. The model was developed in May 2009 in order to quickly foster understanding about the possible dynamics of this new flu variant and to perform rough-cut policy explorations. Later, the model was also used to further develop and illustrate the use of Exploratory System Dynamics models as scenario generators for Exploratory Modelling and Analysis.

The paper starts with an introduction to, and a description of, the exploratory System Dynamics model, followed by a discussion of plausible behaviours, sensitivity, what-if and policy analyses. The model is subsequently used to illustrate the Exploratory System Dynamics Modelling and Analysis approach: base case behaviours are discussed, followed by sensitivity, what-if and policy analyses. Finally some concluding remarks and policy recommendations are formulated.

Keywords: Influenza A(H1N1)v; Deep Uncertainty; Exploratory System Dynamics; Exploratory System Dynamics Modelling and Analysis; Exploratory Modelling and Analysis; Hot Testing/Teaching Cases.

1 Introduction

1.1 Modelling the Influenza A(H1N1)v Pandemic?

In the first days, weeks and months after the first reports about the outbreak of a new flu variant in Mexico and the USA ¹, much remained unknown about the possible dynamics and consequences of the at the time plausible/imminent epidemic/pandemic of the new flu variant, first known as Swine or Mexican flu and known today as Influenza A(H1N1)v. See Table 1, based on the reports of the European Center of Disease Control published between 24 April 2009 and 20 July 2009 for an overview of some of the ‘knowns’ and ‘unknowns’ in the first days, weeks and months after the first signs.

Many nations were at first particularly concerned about the potential loss of human life, and later –after it became clear that the case fatality ratio was moderately low– about the potentially disruptive effects both on health care systems and societies/economies at large in case large fractions of the (active) population would be immobilised simultaneously by the flu.

An Exploratory System Dynamics model was developed shortly after the first signs of a potential outbreak were reported in order to foster understanding about the plausible dynamics of

*Published as: Pruyt E. and Hamarat, C. 2010. The Influenza A(H1N1)v Pandemic: An Exploratory System Dynamics Approach. *Proceedings of the 18th International Conference of the System Dynamics Society*, July 25-29, Seoul, Korea. (Available online at <http://www.systemdynamics.org>.)

¹On 21 April 2009 (Centers for Disease Control and Prevention 2009) and on 30 April 2009 (Trifonov, Khiabanian, Greenbaum, and Rabadan 2009)

Date	24 April	30 April	08 May	20 May	12 June	20 July	21 August
Infectivity	unknown	unknown	unknown	unknown	unknown	unknown	unknown
Ro	unknown	unknown	1–2; prob. 1.4–1.9	1–2; prob. 1.4–1.6	–	–	[R up to 2]
Immunity	unknown	unknown	indications (elderly)	idem	idem	idem	idem
Virulence	unknown	unknown	unknown	unknown	unknown	mild and self-limiting	idem
Incubation period	unknown	unknown	long tail? (up to 8d)	–	median 3–4d range 1–7d	idem	idem
CFR (Mexico)	17%?	–	4%?	2%?	0.4–1.8%?	–	–
CFR (USA)	unknown	unknown	0.1%?	0.1%?	0.2%?	0.4%?	–
CFR (UK)	unknown	unknown	unknown	unknown	unknown	0.3% (–1%)?	0.1–0.2%? (0.35% $\rightarrow$ ex.)
Age distribution	unknown	unknown	older people less affected?	–	skewed tow. younger	idem	idem
Antiviral suscep.	unknown	possible	indications	–	–	–	–
% asymptomatic	unknown	unknown	unknown	unknown	unknown	indications	33–50% (ass.)
Future severity?	unknown	unknown	unknown	unknown	unknown	unknown	unknown
Resp. resources	(ECDC 2009g)	(ECDC 2009f)	(ECDC 2009b)	(ECDC 2009c)	(ECDC 2009a)	(ECDC 2009d)	(ECDC 2009e)

Table 1: Information and unknowns provided by the ECDC from 24 April till August 21

the flu outbreak, and to test policies for dealing with this imminent crisis. The exploratory model presented here is small, simple, high-level, data-poor (no complex/special structures nor detailed data beyond crude guestimates), and history-poor. The model was used in an ex-ante exploratory way: developments were not waited for and uncertainties were amplified and explored instead of reduced or ignored. From a non-exploratory point of view, such simple/aggregated/high-level exploratory models are of course always wrong² –especially with hindsight. We believe nevertheless that the model and method presented here have value, both as an introduction to the A(H1N1)v pandemic and as useful approaches for explorative reflection (ESD) and scenario exploration (Exploratory System Dynamics and Exploratory System Dynamics Modelling and Analysis) in case of imminent outbreaks that are deeply uncertain.

1.2 Exploratory System Dynamics Modelling and Analysis

1.2.1 Traditional versus New Approaches

System Dynamics (SD) is traditionally used for modelling and simulating dynamically complex issues and analysing their resulting non-linear behaviours over time in order to develop and test the effectiveness (and robustness) of structural policies. An important condition for making traditional quantitative SD simulation models –and hence, for traditional SD to be of any use– is that one needs to be able to discern and specify causal relationships and parameter values.

That condition does not necessarily hold under *deep uncertainty*. Deep uncertainty is defined as situations ‘where analysts do not know, or the parties to a decision cannot agree on: (i) the appropriate conceptual models that describe the relationships among the key driving forces that will shape the long-term future [e.g. different drivers and underlying structures than today], (ii) the probability distributions used to represent uncertainty about key variables and parameters in the mathematical representations of these conceptual models, and/or (iii) how to value the desirability of alternative outcomes’ (Lempert, Popper, and Bankes 2003).

Doing more research and making more detailed models may sometimes be helpful for dealing with uncertainty in the sense that such actions could sometimes help to reduce uncertainties related to the specification of causal relationships and parameter values. However, reducing some of the uncertainties will not necessarily make it more predictable: after all, ‘it’s influenza – so it is unpredictable’. And these actions are costly in terms of the time required and may therefore even be problematic in case of imminent or acute crises situations that are characterised by deep

²All models are wrong, but some are just more useful (Sterman 2002).

uncertainty, in other words, in case of insufficient time to do more research and more detailed modelling in order to reduce uncertainties (Pruyt 2010c). Hence, improving models by increasing the level of detail or their size may not help much for dealing with such urgent and uncertain dynamically complex issues. Instead of spending much time trying to develop more detailed models validated on past conditions, aiming to uncover *the* structures underlying the real issue, or predicting future behaviour or *exact probabilities of behaviours*, it may be more useful to:

- quickly explore plausible behaviours with relatively simple non-predictive models;
- explore and analyse the entire space of plausible behaviours that can be generated by large uncertainty ranges for all uncertain parameters and for structurally uncertain model formulations; and
- explore and analyse (the relative effectiveness and) the robustness of policies given all sorts of parametric and structural uncertainties, in other words, in the face of deep uncertainty.

Stand-alone ‘*Exploratory System Dynamics*’ (see subsection 1.2.2) modelling may be useful for the first bullet and thus for rapid threat assessment, and ‘*Exploratory System Dynamics Modelling and Analysis*’ (see subsection 1.2.3) –a combination of Exploratory System Dynamics and Exploratory Modelling and Analysis– may be useful for the second and third bullet and thus for risk assessment to be produced at a later stage, often with more time and information available.

1.2.2 Exploratory System Dynamics

Exploratory System Dynamics (ESD) refers here to the development and use of fast-to-build, relatively-small, highly-simplified, data-poor, easy-to-use SD models to quickly/easily explore possible behaviours and plausible scenarios, and develop a rough idea about the effectiveness (and robustness) of policies. The focus of ESD lies on testing whether behaviours of interest (e.g. trajectories that would require particular attention) are plausible. Scenario analyses, what-if analyses and risk analyses may be useful for ESD in view of exploring behaviours and testing whether behaviours of interest can be generated at all.

The simple ESD models developed are almost certainly wrong, since the issues for which ESD is most appropriate are just too uncertain. Hence, it does not make sense in ESD to derive firm conclusions about specific values or sets of conditions that may lead to ‘this or that behaviour’ in real life, or to interpret outcomes in a probabilistic sense.

Still, ESD models may be much more useful for the purpose of quick and dirty exploration than traditional models. Stand-alone ESD may also be useful as a starting point for deep analysis, of policies for robustly dealing with uncertain and complex issues. However, ESD may not be the best approach for detailed analysis in view of operational implementation or for trying to reduce all uncertainties. More importantly, ESD in isolation may not be sufficiently systematic (for testing the policy robustness) for well-founded decision support for issues characterised by deep uncertainty.

1.2.3 Exploratory Modelling and Analysis

Exploratory Modelling and Analysis (EMA) allows generating insights and understanding about the functioning of systems and the (effectiveness and) robustness of policies, while taking deep uncertainty seriously into account. It consists of using exploratory models –not necessarily SD models– for generating an ensemble of future worlds (tens of thousands of scenarios or more) in order to analyse the space of plausible scenarios and/or to test the effectiveness and robustness of policy options across the ensemble of future worlds. The question for EMA is not *when to measure more* nor *when to model better*, but *how to explore and analyse under deep uncertainty in order to design policies that effectively and robustly improve the system/behaviour given the deep uncertainty*³.

³Readers interested in EMA are referred to (Bankes 1993), (Lempert, Popper, and Bankes 2003), (Lempert, Groves, Popper, and Bankes 2006), (Bankes, Walker, and Kwakkel 2010), (Kwakkel, Walker, and Marchau 2010),

1.2.4 Exploratory System Dynamics Modelling and Analysis

Since EMA requires handy models to generate thousands of plausible scenarios (future states), and ESD requires methods to systematically explore deep uncertainty, they are actually complementary allies (Pruyt 2007) and could be combined into Exploratory System Dynamics Modelling and Analysis (ESDMA). ESDMA consists of (i) developing (one or several) fast-to-build ESD models regarding an uncertain issue and performing a quick and dirty explorative preliminary analysis, (ii) simulating them taking all sorts of uncertainties⁴ into account in order to generate thousands of dynamic behaviours, (iii) analysing the resulting dynamic behaviours given the structural and parametric uncertainties, and/or (iv) testing the (effectiveness and) robustness of policy options given these and other uncertainties. In other words, ESD models are used in ESDMA as EMA scenario generators in order to explore deep uncertainty and support decision-making under deep uncertainty in case of dynamically complex issues. Tools to ease the application of ESDMA are currently under development at Delft University of Technology.

A simple example of ESDMA is provided in section 4 and typical outputs are provided in appendix C. Other recent examples of ESD and ESDMA are discussed in (Pruyt and Hamarat 2010), and (Pruyt 2010b). And (Pruyt 2010c) discusses the inter/national safety and security dimension and crisis aspects of these three examples.

1.3 Goals and Organisation

1.3.1 Goals

The four main goals of this paper are to:

- explain the use of ESD for its own sake, in other words, without combining it with EMA;
- explain and illustrate the use of ESDMA – the combination of ESD and EMA;
- illustrate the application of both ESD and ESDMA to the case of the 2009 flu pandemic;
- and by doing so, shed light on the combination of dynamic complexity and deep uncertainty related to imminent pandemics, in this case the A(H1N1)v pandemic.

1.3.2 Organisation

The basic ESD simulation model is presented in section 2 and the ESD exploration of the behaviour is discussed in section 3. The model is subsequently used for ESDMA in section 4. Concluding remarks and recommendations are formulated in section 5.

SD and SD diagrammatic conventions are briefly explained in appendix A. The stock-flow structure of the automated ESD model is included in appendix B. And appendix C contains detailed ESDMA outputs referred to in section 4.

A ‘hot testing & teaching case’ corresponding to the base model is provided in (Pruyt 2010a). The sim of this base model is available online at <http://forio.com/simulate/simulation/e.pruyt/mexican-flu/>.

Before presenting the ESD model in the next section, it may be useful for readers without SD foreknowledge to get acquainted to SD and its diagrammatic conventions (see appendix A on page 28). Experienced System Dynamicists are advised to continue immediately with the next section.

2 An Exploratory System Dynamics Simulation Model

The modelled world is divided in three regions: the Western World, the densely populated Developing World, and the scarcely populated Developing World. Only the two first regions are included

(Agusdinata 2008), (Bryant and Lempert 2009), et cetera.

⁴Only parametric and some structural/model uncertainties are taken into account in the flu case.

in the model because it is assumed that the scarcely populated regions are causally less important for dynamics of flu pandemics. The basic stock-flow structure of one of the two regional sub-models of the ESD simulation model is displayed in Figure 1. The model is an extremely simple, high-level, symmetric, two-region infection model. Note that parameter and initial values used in subsections 2.1 and 2.2 are/were much more uncertain than the crisp/deterministic description in subsections 2.1 and 2.2 hints at. These crisp parameter and initial values will be varied/explored in what follows.




Figure 1: Stock-flow diagram of region 1 (the ‘Western World’) of the ESD two-region flu model

2.1 Region 1: The ‘Western World’

Infections make that people from the *susceptible population* get the flu and migrate from the *susceptible population* to the *infected population* (initially equal to 10 persons). The number of *infections* equals the product of the *susceptible population*, the *contact rate*, the *infection ratio*, and the *infected fraction*. The *infected fraction* equals the *infected population* divided by the *total population*. The average *infection ratio* of this flu variant in the Western World was –as a start– set to 10% infections per (close) contact [which may on the one hand be rather high]. It was assumed that the *normal contact rate* in de Western World amounted to 50 (close enough) contacts per person per month [which may on the other hand be rather low]. The *contact rate* is then the product of the *normal contact rate* and the *impact of the infected population on the contact rate*. The latter variable is included in order to take the contact rate reduction due to public anxiety and illness into account, with following function: $1 - (\text{infected fraction region 1})^{(1/\text{root contact rate region 1})}$, with the *root contact rate of region 1* initially equal to 1.5.

Members of the *infected population* flow –after an average *recovery time* of 2 weeks⁵– to the *recovered population* (which is initially empty), that is, if they recover. If not, they flow to the *deceased population* stock that can be attributed to the new flu variant. How many people recover or die depends on the *fatality ratio*, initially set at 0.1%.

⁵It is implicitly assumed here that incubation and illness are equally infectious, and together take up to 2 weeks.

At the first signs of the outbreak, it was not clear whether there would be a seasonally and/or permanently *immune population*. Assuming the existence of seasonal and permanent immunity, the *susceptible population* in the Western World was initially set to 330.000.000 persons and the *immune population* to 270.000.000 persons. An additional structure was added in order to model the seasonal dynamics between both populations: the *susceptible population* grows towards the winter season, and the *immune population* grows towards the summer season.

The flow between these populations –the net *susceptible to immune population flow*– equals the gap between the *normal immune population* and the *immune population* divided by the *susceptible to immune population delay time*. However, this flow cannot be greater than the *susceptible population* divided by the delay if there is a net flow towards the immune population, and the return flow cannot be greater than the *immune population* divided by the delay if there is a net flow towards the susceptible population. This delay was set to 1 month.

The *normal immune population* equals the product of the *total population* and the *normal immune population fraction*. This *normal immune population fraction* of the Western World was modelled as a sinus function fluctuating between 30% in month 0, to 70% in month 6, to 30% in month 12, and so on. In other words, it was assumed in the base model that 30% of the population is permanently immune, and an additional maximum of 40% could be seasonally immune in summer.

2.2 Region 2: The Densely Populated Developing World

The flu epidemic in the Western World may be strongly influenced by the development of the flu in densely populated Developing Countries. A similar submodel for a second region, namely the densely populated part of the Developing World was therefore added.

There are some small differences between the two regions: it was assumed that the average *normal contact rate* in this second region amounts to 100 (close) contacts per person per month [which may be rather low too]. The *infection ratio* of the Developing World was assumed to be equal or slightly higher, e.g. 15% or 0.15 infections per contact [which may be rather high too]. It was assumed that the *susceptible population* of this region initially amounted to 2.000.000.000 persons, the *immune population* to 1.000.000.000 persons, and the *infected population* to 100 persons. The majority of these countries are located in the Southern hemisphere, and it is assumed that their populations have lower levels of immunity: the *normal immune population fraction* was assumed to vary according to a sinus function from 30% in month 0, to 10% in month 6, back to 30% in month 12, et cetera. In other words, it was assumed that 10% of the population is permanently immune, and a maximum of 20% of the population could be seasonally immune (in the northern hemisphere summer).

There is of course also contact between these regions (between the Western World and the densely populated Developing World). In the base case, it was assumed that the *interregional contact rate* amounts to 0.1 (close) contacts per person per month. An additional term was therefore added to the *infections* variables of both regions.

2.3 Possible Shortcomings and Extensions

The model presented above is rather aggregated and very simple – even simplistic. The model could be extended easily with asymptomatic cases, an incubation period, gradual loss of immunity of (part of) the recovered population and different infectivity ratios. These changes would turn the Stock-Flow structure into a structure as in (Pruyt 2009). The structure adding seasonal and permanent immunity described above may also be too simple (and too optimistic and therefore slightly inappropriate for simulating worst credible scenarios). A major simplification of the model relates to the implicitly assumed homogeneity of the different populations. The assumed homogeneity is certainly wrong: risk and age groups matter in epidemiology. This faulty assumption needs to be abandoned if the purpose of the modelling endeavor is answering detailed questions

—such as which (risk or age) groups to target— or replicating outbreak after the facts, instead of exploring plausible future outbreaks given the uncertainties remaining. However, freeing the model from this faulty assumption would require much more detailed modelling (e.g. of age and risk groups, divided in healthy and unhealthy subgroups, maybe even subdivided according to their underlying conditions); of social and geographic network structures; of different countries and/or different populated areas, et cetera), hence making the models bigger and less manageable, pushing the analysis further away from the exploratory objective.

These extensions were not included given the ‘known knowns’ at the time of modelling and our specific aims for the ESD (a quick and dirty analysis of plausible/worst credible scenarios and testing the effectiveness of policies for these scenarios) and the ESDMA (an exploration of the uncertainty space of plausible worst credible scenarios and the robustness of policies). It would nevertheless be interesting to refine the model —now that the urgency is practically nil— with our current knowledge about its characteristics and detailed social-geographic information.

3 ESD: Quick and Dirty Exploration of Uncertain Behaviours and Policy Effectiveness

3.1 ESD Base case behaviours

Before discussing any ‘base case’ behaviours, it should be noted that the behaviours shown in the current subsection cannot be interpreted as ‘most likely’ behaviours. These four base case scenarios are just used here to show that this model generates —given the set of values introduced in section 2— rather different dynamic behaviours for different values of the *infection ratios*. BC1 combines an *infection ratio of region 1* of 10% and *infection ratio of region 2* of 10%; BC2 of 15% and 15%; BC3 of 20% and 20%; and BC4 of 10% and 15% [again, these values are rather high, but are somewhat compensated for by the rather low contact rates].

The graphs in Figures 2 and 3 give some idea about the model behaviour of the structure displayed in Figure 1. In the Western World, seasonal immunity shifts a larger fraction of the (Western) population from the *immune population* to the *susceptible population* and back. Depending on the base case scenario, *infections* drain the entire *susceptible population* except part of the seasonally immune. The number of *infections* —and hence the *infected fraction of the population*— behave rather differently for these 4 base cases. The *impact of the infected population on the contact rate* (with a root contact rate of 1.5) somewhat decelerates the acceleration of the outbreaks. The recoveries and deaths show the same pattern of behaviour (namely the cumulative of the first order delay of the *infections*).

Figure 3 shows that, in the developing world, the small immunity fraction leads —in combination with slightly higher infection ratios than in the Western World— to a catastrophic outbreak, with a very high infected fraction, rapidly draining the *susceptible population*.

3.2 ESD Sensitivity Analyses

The sensitivity of the model to changes in the variables is tested in this subsection by means of univariate sensitivity analyses and multivariate sensitivity analyses (simultaneously changing uncertain variables) taking relatively broad uncertainty ranges into account.

3.2.1 Parametric Uncertainty

Table 2 lists uncertain model variables and the distributions used in this and the next subsection to take this uncertainty into account. All univariate sensitivity analyses below start from BC4, with infection ratios of 10% and 15%. The graphs in Figures 4 and 5 give an indication of the sensitivity of the model in terms of the *infected fraction* —as a proxy for societal disruption— and the *deceased population* for the Western World. The graphs do not show the results of classic




Figure 2: Base case behaviours of the exploratory flu model for region 1




Figure 3: Base case behaviours of the exploratory flu model for region 2

sensitivity analyses, but an exploratory variant where 20 values are sampled (Latin Hypercube) from these distributions.

The graphs in Figures 4(a) and 4(b) show that the model is somewhat sensitive to changes in the *additional seasonal immune population fractions*. The graphs in Figures 4(c) and 4(d) show that the model itself is almost not sensitive to changes in the *fatality ratios*, but that the values of the *deceased population* output indicator are. The graphs in Figures 4(e) and 4(f) show that the model is sensitive, both in terms of patterns and strengths (bigger first and/or second peak and less/more deaths), to changes in the normal interregional contact rate. The graphs in Figures 4(g) and 4(h) show that the model is very sensitive to changes in the *permanent immune population fractions*: different values of the *permanent immune population fractions* lead to very different values for the *deceased populations*, mainly caused by different strengths of the second peak. Hence, this variable may be critical for monitoring and (adaptive) policymaking.

The graphs in Figures 5(a) and 5(b) show that the model is very sensitive to changes in the *recovery times*: different *recovery times* lead to very different values for the *deceased populations*, mainly caused by different strengths of the second peak. This variable may also be a critical variable to be monitored and used for adaptive policy making. And the graphs in Figures 5(c) and 5(d) show that the model is somewhat more sensitive to changes in the *root contact rates*. The graphs in Figures 5(e) and 5(f) show that the model is behaviourally sensitive to changes in the *infection ratio*. The graphs in Figures 5(g) and 5(h) show that model is also behaviourally sensitive to changes in the *normal contact rate*, but –starting from BC4– only for a wave in the northern hemisphere winter, or later. Starting from the BC4 parameter value set, the output indicators are only slightly numerically sensitive to changes in the *susceptible to immune population delay time* and negligibly sensitive to changes in the *initial immune fractions of the populations*. It should be noted though that these in/sensitivities are ‘local’: they only apply to the base model presented in section 2 in combination with the BC4 parameter set.




Figure 4: Univariate sensitivity analyses (1/2)




Figure 5: Univariate sensitivity analyses (2/2)

variable:	UVSA	distribution	low	high
add. seasonal immune population fraction region 1	1	random uniform	0	0.5
add. seasonal immune population fraction region 2	1	random uniform	0	0.5
fatality ratio region 1	2	random uniform	0.0001	0.01
fatality ratio region 2	2	random uniform	0.0001	0.1
initial immune fraction of the population of region 1	3	random uniform	0	0.5
initial immune fraction of the population of region 2	3	random uniform	0	0.5
normal interregional contact rate	4	random uniform	0.01	1
permanent immune population fraction R1	5	random uniform	0	0.5
permanent immune population fraction R2	5	random uniform	0	0.5
recovery time region 1	6	random uniform	0.1	0.75
recovery time region 2	6	random uniform	0.1	0.75
susceptible to immune population delay time region 1	7	random uniform	0.5	2
susceptible to immune population delay time region 2	7	random uniform	0.5	2
root contact rate region 1	8	random uniform	0.01	5
root contact rate region 2	8	random uniform	0.01	5
infection ratio region 2	9	random uniform	0	0.15
infection ratio region 1	9	random uniform	0	0.15
normal contact rate region 2	10	random uniform	10	100
normal contact rate region 1	10	random uniform	10	70

Table 2: Uncertain variables and the distributions, lower bounds and higher bounds used – note that the upper bounds for the fatality ratios and contact rates are lower than in section 4

When changing the uncertain parameters simultaneously (using Latin Hypercube sampling), the 1000 individual traces displayed in Figure 6(a) show that there are many different plausible behaviours given this model and the uncertainties considered: from plausible –but rather unlikely if we can talk about likelihoods– catastrophic (mostly early) outbreaks, over considerable multiple peak epidemics, to many rather mild epidemics. Although it is difficult to see in Figure 6(a) –because attention is naturally drawn towards the enveloping contours of the ensemble of scenarios– it should be noted that in most cases the pandemic is *not catastrophic*: Figure 6(b) shows that most values for the infected fractions remain rather low. Although there are strong philosophical arguments⁶ and practical arguments⁷ against interpreting these outcomes as ‘confidence’ bounds or probabilities, they are displayed here to show that the large majority of simulated evolutions of the are mild at worst. Most simulated flu waves are not more deadly or disruptive to society than normal flu waves and cannot be classified as catastrophic, which makes them less interesting from a ‘worst credible scenario’ point of view: here we are interested in the entire space of *plausible* scenarios but especially in the worst credible ones, even though they are extremely unlikely.

Such risk analyses could of course also be used to explore the overall impact of important assumptions (see subsection 3.2.2) and of promising policies (see section 3.3).

3.2.2 Structural/Model Uncertainty

Different model formulations have been tested as part of a ‘Structure-Oriented What-If Analysis’. The model used for this analysis is displayed in Figure 16 on page 30. This version of the model can be manipulated easily for exploratory purposes: submodels of this version of the model can easily be turned on/off for simulating the model with one region or two regions, with/out seasonal and/or permanent immunity, with/out flu mortality, etc. In other words: structurally different simulation models can be simulated by switching structures on/off.

Figures 7(a) and 7(b) show the MVSA Western World outputs without immunity (immunity

⁶EMA and confidence bounds are philosophically incompatible.

⁷The confidence bounds displayed here are related to the values at each moment in time (vertical differentiation), but not to the patterns (horizontal differentiation) which would be more interesting.





Figure 6: MVSA of the flu model (with different degrees of seasonal and permanent immunity)



(a) 1000 individual traces of the MVSA base case without immunity (b) ‘Confidence bounds’ of 1000 runs without immunity

Figure 7: A structural what-if MVSA of the flu model without immunity

being switched off). Then, a rapid outbreak of increased strength seems –not surprisingly– more plausible. Again, it should be noted that most runs are not catastrophic and that attention is naturally drawn towards the enveloping contours of the ensemble of runs. The other MVSA’s – using the other submodel switches– are not displayed here because their results are less conclusive and/or interesting.

3.3 ESD Policy Analyses – from Static to Adaptive

After analysing plausible dynamic behaviours, we will now test the effectiveness (and robustness) of policies for dealing with those behaviours, especially worst credible ones. Static policies are tested in subsection 3.3.1 and dynamic policies are tested in subsection 3.3.2. Note however that one of the major known unknowns –the effectiveness of policies– is not explored in depth here.

3.3.1 Static Policy Analysis

Figure 8 displays the effect of two policies and a combined policy applied to the base case MVSA:

- Policy 1: a 100% effective vaccine becomes available at the start of month 10 (about 5 months after the first signs of a new flu strain) and 70% of the population living in developed countries is vaccinated (instantly);
- Policy 2: a contact rate reduction of 50% is achieved through a plethora of social measures (closure of schools, bans on mass meetings, etc.) from month 4 till month 10;

- Policy 3 = policy 1 + policy 2: a contact rate reduction of 50% is achieved through all sorts of social measures from month 4 to month 10 and a vaccine becomes available at the start of month 10, and 70% of the population living in developed countries is vaccinated (instantly and 100 effectively).

Figure 8 shows –in this model and for these uncertainty ranges– that:

- policy 1 effectively deals with catastrophic epidemics after month 10 but not before;
- policy 2 mostly reduces the risk of a catastrophic epidemic in the first months, shifting peak outbreaks to the following Northern hemisphere winters;
- the combination of these static policies effectively reduces most of the plausible peak outbreaks.

3.3.2 Dynamic/Adaptive Policy Analysis

The policies simulated in the previous subsection are static in the sense that critical indicators are not monitored to adapt pre-specified policies given the information that (gradually) becomes available. But in case of an imminent flu pandemic, monitoring *is* very important, because:

- the development of vaccines may take more time than is available to avoid a first plausible wave (in this case a 2009 northern hemisphere summer wave);
- additional information will undoubtedly become available during the first few months which could be used to base the number of vaccines ordered on;
- social/societal crisis measures could be taken too if critical thresholds are crossed.

Following critical variables –to which the model is rather sensitive– may be really worth monitoring during the first months: the fatality ratio, the recovery time, the infection ratio, and the infected fraction. The additional information gathered may help to narrow down the space of plausible evolutions and eliminate most of the (previously) plausible scenarios and policies, which is very useful, both for supporting decision making and crisis management.

Moreover, there are some problems with the static policies presented in subsection 3.3.1:

- Static ex-ante vaccination policies –vaccinating a particular fraction of the population irrespective of (gradually gathered information about) the infectivity, the fatality, permanent and seasonal immunity, et cetera– are almost always suboptimal and lead to overinvestments (e.g. 100% coverage in case of a mild flu with no fatalities), underinvestments (e.g. 0% coverage in case of a very lethal variant), and only sometimes to exactly the right policy.
- The static social measure simulated above –a 50% contact rate reduction over a period of 6 months no matter what the outbreak looks like– is always worse –in terms of societal disruption– than the social disruption of the flu in the absence of social measures. [It should be noted that this argument does not hold if the fatality ratio is high(er than normal) and/or the hospitalisation rate is unusually high.]

Policies should indeed depend on the situation, and from a dynamic perspective, on the (expected) evolution of the situation: static robust policies are just not good enough. And although real-world decision making processes are mostly iterative and often lead to adaptive decision making, new information gathered through monitoring of critical indicators is not always exploited to the full extent –especially not under deep uncertainty– and is almost always ad-hoc, mostly reactive, and often used when it is already too late.

Dynamic/adaptive policies are proactively designed to handle (unpredictable) risks by monitoring and prespecified responses. In this subsection we will test dynamic versions of the policies tested in subsection 3.3.1. The static policies are adapted as follows:




Figure 8: Quick and dirty ESD policy analysis: *infected fraction region 1* for 3 **static** policies

- Policy 1: For the adaptive version of policy 1 it is assumed that after the first signs of an outbreak and an initial basic order of vaccines there is some time to monitor the fatality ratio (and or other aspects such as combination of infectivity, immunity, et cetera) before final amounts of vaccines need to be ordered. The following (illustrative) adaptive policy has therefore been added to the model: if the fatality ratio in region 1 is –after 3 months– equal or greater than 1%, then vaccines will be ordered for 60% of the population, if it is greater than 0.1%, then vaccines will be ordered for 50% of the population, if it is greater than 0.01%, then vaccines will be ordered for 40% of the population, and else no vaccines will be ordered at all. Although most likely still expensive, this policy will for each variant be cheaper than the static policy (always 70% coverage). Here too, it is assumed that the vaccine becomes available at the start of month 10 (about 5 months after the first signs).
- Policy 2: A new variable –the *orchestrated contact rate reduction*– is added to the model in order to make policy 2 adaptive. It is an S-shaped function of the *infected fraction of region 1* with couples (0,0), (0.05,0.05), (0.1,0.2), (0.15,0.5), (0.2,0.75), (0.3,0.85), (0.4,0.9), (1,1). In other words, the *orchestrated contact rate reduction* first increases exponentially to 50% in case the *infected fraction in region 1* increases to 15% and then converges to 100% if the *infected fraction in region 1* increases to 100% (which is extremely unlikely to happen). The *contact rate of region 1* then equals the product of the *normal contact rate of region 1*, the *impact of the infected population on the contact rate of region 1*, and 100% minus the *orchestrated contact rate reduction*.
- Policy 3 is a combination of policy 1 and policy 2. Policy 2 remains activated, also after the start of vaccination campaigns.

Figure 9 shows that the dynamic version of policy 1 could be used for post-summer wave control, that the dynamic policy 2 could be used for ‘peak shaving’, and that policy 3 could be used for ‘peak shaving’ and post-summer wave control, in all cases at a much lower (societal and financial) cost than the static policies.

4 ESDMA: Systematic Exploration of Uncertain Behaviours and Policy Robustness

The ESD approach is an interesting quick and dirty assessment approach of plausible dynamics and policy effectiveness. However, the analyses in the previous section may not be deep and broad enough from a ‘deep uncertainty perspective’, because:

- the uncertainty ranges considered are too narrow and conservative (broader ranges would be more difficult to handle by ESD alone),
- the analysis of the dynamic behaviours is too shallow (but deeper analysis is more difficult to handle and time consuming by means of ESD in isolation),
- and the robustness is not tested over the entire uncertainty space (ESDMA is more appropriate for that although it could be achieved by ESD alone).

However, the work presented in sections 2 and 3 form –from the point of view of ESDMA– a good starting point for further exploration of uncertainties and policies. That is why we turn to it in the current section.

The same structure as in section 3 is used here: first, simulation runs are analysed (step 3 of ESDMA), then the robustness of potential policies –first static, then dynamic– is tested.




Figure 9: Quick and dirty ESD policy analysis of 3 **dynamic** policies

4.1 ESDMA: Analysis of Behaviours

In this subsection, we will actually perform (more or less) the same Exploratory Analysis twice, each time with a different sampling technique. Both are informative and may generate interesting insights. Applying both helps to investigate and illustrate the trade-off between the degree of detail/structuredness and the degree of inclusiveness:

The Full Factorial Design (FFD) discussed in subsection 4.1.1 can be used for detailed/structured exploration of the entire uncertainty space defined by a few uncertainties. Only a few uncertainties can be explored simultaneously because the FFD is computationally extremely expensive.

Latin Hypercube Sampling (LHS) discussed in subsection 4.1.2 is computationally less expensive and allows to include more uncertainties in the analysis, but with a loss of detail/structuredness. LHS could therefore be used to combine more uncertainties, but in a less detailed/structured way.

4.1.1 Full Factorial Design: Few Uncertainties but Full Coverage

This model contains 19 uncertain parameters which could be explored. However, covering all these parameters with a Full Factorial Design (FFD) for only 2 values for each of these parameters already means 2^{19} –or almost half a million– simulation runs. We therefore selected –based on the analyses in section 3.2– only 7 uncertain parameters for which we included the uncertainty ranges displayed in Table 3. This table shows that the parameter ranges used in this subsection are rather high and broad –higher and broader than the ranges in the previous section. But it should be noted that many uncertainties –such as the those surrounding the fatality ratios, fixed here at 0.1%– are not explored.

Parameter	Value 1	Value 2	Value 3	Value 4	Value 5	# of different values
Infection Rate Region 1	0.05	0.10	0.15	0.20	0.25	5
Infection Rate Region 2	0.05	0.10	0.15	0.20	0.25	5
Normal Contact Rate Region 1	20	40	60	80	100	5
Normal Contact Rate Region 2	20	40	60	80	100	5
Recovery Time Region 1 (in months)	0.2	0.4	0.6	0.8	1.0	5
Recovery Time Region 2 (in months)	0.2	0.4	0.6	0.8	1.0	5
Normal Interreg. Contact Rate	0.2	0.4	0.6			3

Table 3: Parameter values tested in the FFD (with a fixed case fatality ratio of 0.1%)

The reason for using these extreme rather ranges here is that we want to explore possible catastrophes beyond the worst past catastrophes and beyond our greatest fears for the future in order to see what it could plausibly lead to and whether and how we would still deal with them.

Figure C1 in the appendix shows 1875 runs per graph for infection ratios of 5% to 25% –which could be considered ‘worst plausible’ but less realistic, and Figure 10 zooms in on 1875 runs per graph of *infected fractions* of 5% to 10% –which may be considered more ‘credible’ and realistic by many. Figure C2 shows –by zooming in on the 10%–10% plot– that few extreme waves are much more visible than the much larger number of less catastrophic waves. Figure 10 shows that a considerable infection ratio is needed for really causing a serious epidemic peak in the first northern hemisphere summer (by which time vaccines would not be available), and that this peak could already be rather catastrophic in some of the cases (but these cases are rather mild compared to other cases displayed in Figure C1). And the first peaks of many of the worst plausible cases are –as could be expected– also much more catastrophic. Although not pursued here, these trajectories could be further analysed with data analysis/mining tools.

4.1.2 Latin Hypercube Sampling: More Uncertainties but Reduced Coverage

Although the detailed analyses in the previous subsection give a good visual impression of plausible dynamics, it only covers part of the uncertainties and is therefore insufficient from the perspective




Figure 10: Infected fraction for a (relatively mild and rather credible) subset of the ensemble of plausible scenarios displayed in Figure C1

of exploration/analysis of deep uncertainty. Although it was good to be able to track the trajectories in a structured and detailed way, we will now move to a more inclusive and aggregated perspective: since FFD is not feasible for 19 uncertainties, a more efficient sampling technique –Latin Hypercube Sampling (LHS)– is used here. More uncertainties are included, and broad(er) upper and lower bounds are used here (see Table 4) and in the remainder of the paper, except for four parameters:

- The two infection ratios are limited to maximum 10% – which corresponds to the ‘credible’ range considered in Figure 10.
- The two *susceptible to immune population delay time* parameters have been excluded from the analysis for both regions because previous analyses showed that they do not substantially influence the model outcomes.

We already mentioned that there are at least three aspects worth monitoring: the (fraction of, or number of) fatalities, the societal disruption (or a proxy such as the maximum infected fraction of the population), and the time available before the worst wave (or a proxy such as the time before the peak infected fraction). Combined with information about the time needed for appropriate action (e.g. development and production of an appropriate amount of vaccines) they

Parameter	Lower Limit	Upper Limit
additional seasonal immune population fraction region 1	0.0	0.5
additional seasonal immune population fraction region 2	0.0	0.5
fatality ratio region 1	0.0001	0.1
fatality ratio region 2	0.0001	0.1
initial immune fraction of the population of region 1	0.0	0.5
initial immune fraction of the population of region 2	0.0	0.5
normal interregional contact rate	0.0	0.9
permanent immune population fraction region 1	0.0	0.5
permanent immune population fraction region 2	0.0	0.5
recovery time region 1	0.2	0.8
recovery time region 2	0.2	0.8
root contact rate region 1	1.0	10.0
root contact rate region 2	1.0	10.0
infection ratio region 1	0.0	0.1
infection ratio region 2	0.0	0.1
normal contact rate region 1	10	200
normal contact rate region 2	10	200

Table 4: Parameter ranges for the LHS with parameterised fatality ratios (0.01% – 10%) and reduced –more ‘credible’– ranges for the infection ratios (0% – 10%)

determine the need for and urgency of taking action.

Figure 11 shows the 3D scatter plot of a LHS for 20000 different samples with 17 uncertain parameters in terms of these three aspects (the total cumulative number of fatalities, the maximum infected fraction and the time index of the maximum infected fraction –all for the Western World). [The corresponding projections are displayed in figure C4 in the appendix.]

The worst catastrophic peaks (in terms of urgency-fatality-disruptiveness) seem to occur in the first few months, especially in the northern hemisphere summer and to a lesser extent in the following northern hemisphere winter. These early catastrophic cases are all characterised by high infection ratios and/or long recovery times – which indeed cause acute outbreaks with large infected fractions. It could be concluded that variants with a very severe first summer peak and variants with a severe winter peak in the first winter after the initial outbreak may –without any policies/measures in place– be catastrophic. However, if vaccines can be developed and produced at a large scale before the winter peak, then the second (northern hemisphere winter) peak could possibly be contained, at least to some extent. Hence, the need for other (e.g. social distancing and other) measures to be able to deal with plausible catastrophic waves before vaccines are available and administered. Without large-scale vaccination campaign, this flu may afterwards still lead to flu waves in subsequent winters.

Although a probabilistic interpretation of these results would not be methodologically sound, it is interesting to note that 43.7% of the runs (or 8749 out of 20.000 runs) without vaccination and social distancing policies result in more than 600.000 simulated fatalities or 0.1% of the region 1 population, 26.9% of the runs result in more 1% of the region 1 population (or 5381 out of 20.000 runs), and 13.2% in more than 2.5% of the region 1 population (or 2631 out of 20.000 runs). A catastrophic pandemic was therefore, given this model and all the unknowns, not implausible. Higher upper bounds for the fatality ratios and infection ratios (although not displayed here) would lead to even more and more severe early catastrophic outbreaks (first and second waves).

The big picture conclusion from these LHS simulations is that catastrophic outbreaks are indeed plausible, especially acute/imminent ones (i.e. with first or second wave peaks), and that policies are needed for dealing with the catastrophic cases. In the next section, we will investigate whether static and/or dynamic policies could be designed for keeping the catastrophic cases under control without too much disruption.




Figure 11: Scatter plot with projections of the LHS 20000. X-axis: 0–48 months; Y-axis: 0–50% infected fraction; Z-axis: 0–50.000.000 fatal cases (see also Figure C3)

4.2 Robust Policy Making

Exploratory analysis may be helpful, but testing the robustness of potential policy options under deep uncertainty may be even more important. That is what is illustrated in this subsection.

4.2.1 Static Policies

Figure C4 shows the effects of different values for the static versions of the vaccination policy (Policy 1). The results are displayed for two different variants: a best case ‘maximally effective vaccination programme’ (top row) and a worst case ‘minimally effective vaccination programme’ (bottom row). The vaccination programme is maximally effective if only those people that are *not* immune are vaccinated. The vaccination programme is minimally effective if those people that *are* immune are vaccinated first. Reality will be somewhere between those two variants, since for an imminent outbreak of a new flu variant it is (rather to very) uncertain who is and who is not (seasonally or permanently) immune. Vaccinating 40% of the population in the best case seems to be about as good as vaccinating 70% of the population in the worst case (both resulting – as indicated between square brackets – in about 32% of the runs with more than 600.000 fatalities on the region 1 population of 600.000.000 – down from 43.7% without policy). Given the uncertainty and the need to prepare for the worst case, it seems as though 60–70% of the population may have to be vaccinated. But even 100% coverage does not prevent early catastrophic outbreaks – the ones that break out before full implementation of the vaccination campaign. Hence, additional

measures are needed to deal with the pre-vaccination period.

Figure C5 shows the effects of different values for the static version of Policy 2 (permanent social distancing measures that result in a permanent reduction of the contact rate) statically implemented over the entire time horizon (just to show what different percentages of contact rate reduction could plausibly lead to). This figure shows that this measure is effective in some of the cases, but not in all. The static contact rate reduction needs to be rather high for a significantly robust reduction of the plausible number of early catastrophic waves (the percentages of runs with a population fatality ratio of more than 0.1%, 1%, and 2.5% are displayed between square brackets). But high static contact rate reductions are also disruptive from a societal point of view, maybe even more disruptive than the immobilisation caused by the flu. Hence, it could be argued that such drastic measures would only be sound for exceptionally infective and/or lethal outbreaks (see subsection 4.2.2).

Figure C6 shows the effects of different values for the combination of the static versions of Policy 1 (more precisely the worst case ‘minimally effective vaccination programme’) and Policy 2 (applied here only until November 2009, because it is assumed that by then the vaccination campaign has been implemented). The values between square brackets (again the % of 20.000 runs with more than 600.000 fatalities on the region 1 population of 600.000.000) show that it pays to invest in more and more vaccination (note the increasing returns to scale), more so than forcing more and more social distancing. However, applying *both* is even better: combining a vaccination coverage of 70% with a contact rate reduction of 50% leads to a reduction of the percentage of runs with more than 0.1% fatalities in region 1 from 43.7% down to 19.5% (see Table 5), to a reduction of the percentage of runs with more than 1% fatalities in region 1 from 26.9% down to 6% (see Table 6), and to a reduction of the percentage of runs with more than 2.5% fatalities in region 1 from 13.2% down to 1.6% (see Table 7).

No pol: 43.7%	Policy 1 =30%	Policy 1 =40%	Policy 1 =50%	Policy 1 =60%	Policy 1 =70%
Policy 2=20%	41.0%	38.8%	35.2%	30.5%	26.5%
Policy 2=30%	40.5%	38.2%	34.3%	29.0%	23.7%
Policy 2=40%	40.1%	37.8%	33.5%	27.8%	21.5%
Policy 2=50%	39.7%	37.3%	33.0%	26.9%	19.5%
Policy 2=60%	39.3%	36.8%	32.4%	26.3%	18.2%

Table 5: Fraction of runs with a total population fatality ratio (PFR) greater than 0.1%

No pol: 26.9%	Policy 1 =30%	Policy 1 =40%	Policy 1 =50%	Policy 1 =60%	Policy 1 =70%
Policy 2=20%	24.3%	21.7%	18.3%	14.4%	12.2%
Policy 2=30%	24.0%	21.4%	17.6%	12.9%	10.0%
Policy 2=40%	23.9%	21.2%	17.1%	11.7%	7.7%
Policy 2=50%	23.8%	21.1%	16.8%	10.9%	6.0%
Policy 2=60%	23.7%	20.9%	16.6%	10.5%	4.9%

Table 6: Fraction of runs with a total population fatality ratio (PFR) greater than 1%

No pol: 13.7%	Policy 1 =30%	Policy 1 =40%	Policy 1 =50%	Policy 1 =60%	Policy 1 =70%
Policy 2=20%	10.9%	8.9%	6.8%	5.6%	5.4%
Policy 2=30%	10.8%	8.5%	5.9%	4.3%	3.9%
Policy 2=40%	10.7%	8.3%	5.3%	3.2%	2.6%
Policy 2=50%	10.7%	8.3%	5.0%	2.3%	1.6%
Policy 2=60%	10.7%	8.2%	4.9%	1.9%	0.7%

Table 7: Fraction of runs with a total population fatality ratio (PFR) greater than 2.5%

Note that these static social distancing policies are prohibitively expensive in terms of societal disruption. That is why many countries actually already have some sort of dynamic/adaptive

social distancing policy⁸. Adaptive vaccination policies may be needed too.

4.2.2 Adaptive Policies

The adaptive vaccination policy (Policy 1) simulated here for illustrative purposes is defined in function of the *population fatality ratio*. The vaccination coverage is at least 20% –assumed to be the minimum to start up vaccine development and production. It is assumed that the coverage finally planned/ordered for will be based on information regarding the observed population fatality ratio according to pre-specified values. Here we test a rather low-coverage variant: the coverage aimed for is 60% in case of a total population fatality ratio $>2.5\%$, else the coverage is 20%. These values result –for our upwardly-biased sample set– in an average vaccination coverage of 50%⁹.

Figure 12(b) and the second row in Table 8 illustrate that this adaptive policy appropriately deals with the outbreaks after the end of the vaccination campaign, but not before.

The first adaptive social distancing policy tested (Policy 2a) is based on monitoring of the infected fraction according to the function in Figure 12(c). Figure 12(e) and the third row in Table 8 illustrate that this adaptive policy seems to halve the worst infected fractions, but not so much –or almost not– the cumulative number of fatalities. Combining this variant of Policy 2 with Policy 1 only really shows improvement in terms of cumulative fatalities after the vaccination campaign and the infected fraction in all cases (see Figure 12(g) and the fifth row in Table 8).

Hence, a second adaptive social distancing policy was tested (Policy 2b) based on monitoring of the fatality ratio and the infected fraction according to the function in Figure 12(d). This function transforms the product of the fatality ratio and the actual infected fraction into social distancing actions. Figure 12(f) and the fourth row in Table 8 illustrate that this adaptive policy –although still simplistic– reduces the infected fraction as well as the cumulative number of fatalities. Combining this variant of Policy 2 with Policy 1 improves the outcomes before and after the infection campaign (see Figure 12(h) and the sixth row in Table 8).

	% of runs PFR $>0.1\%$	% of runs PFR $>1\%$	% of runs PFR $>2.5\%$	# of runs PFR $>0.1\%$	# of runs PFR $>1\%$	# of runs PFR $>2.5\%$
Without policies	43.7%	26.9%	13.2%	[8749]	[5381]	[2631]
Dyn. Policy 1	36.7%	18.8%	8.3%	[7343]	[3766]	[1672]
Dyn. Policy 2a	43.6%	26.7%	12.8%	[8723]	[5349]	[2569]
Dyn. Policy 2b	41.9%	24.8%	11.0%	[8384]	[4968]	[2198]
Dyn. Policy 3a	36.2%	18.2%	7.7%	[7236]	[3646]	[1544]
Dyn. Policy 3b	34.9%	16.1%	4.4%	[6974]	[3221]	[870]


Table 8: Percentage and number of runs with population fatality ratios $> 0.1\%$, 1% and 2.5%

Both combined policies seem to be effective, but not robust, in the sense that there are still catastrophic outbreaks. Stronger (vaccine coverage) policies perform much better. And adding adaptive anti-viral policies may help even further to reduce the cumulative number of fatalities.


However, the point we really wanted to make here is that these relatively light versions of the adaptive policies perform better than the corresponding static policy while at the same time being significantly less disruptive in societal terms (only social distancing measures when and proportional to what is needed at any moment in time).

⁸See for example www.flu.gov/professional/community/commitigation.html.


⁹There are 15015 runs in this sample set with a total population fatality ratio $>2.5\%$, 3003 runs with a total population fatality ratio $>1\%$, 1801 runs with a total population fatality ratio $>0.1\%$, and 181 runs with a total population fatality ratio $<0.1\%$. The reason for this upward bias is the uniform sampling of the fatality ratio between 0.01% and 10%




(a) No Policy




(b) Adaptive Policy 1: vaccination based on the fatality ratio




(c) Adaptive Policy 2a: social distancing in function of infected fraction




(d) Adaptive Policy 2b: social distancing in function of fatality ratio




(e) Adaptive Policy 2a: social distancing based on monitoring of infected fraction



(f) Adaptive Policy 2b: social distancing based on monitoring of fatality ratio



(g) Policy 3a: Policy 1 x Policy 2a



(h) Policy 3b: Policy 1 x Policy 2b

Figure 12: Effect of simple adaptive policies starting from and compared with the base case LHS

5 Conclusions

5.1 Conclusions regarding the 2009 and future flu pandemics/epidemics

The 2009 A(H1N1)v flu epidemic/pandemic received a lot of (media) attention, almost creating a panic, until it became clear that its fatality ratio was not higher than the yearly flu. However, the 2009 could equally well have been a catastrophic outbreak. But even then, it makes sense to perform a rapid but decent threat assessment, including a quick exploration of plausible dynamics and a preliminary assessment of the effectiveness and robustness of potential policies under deep uncertainty. Newly available information –obtained through focussed monitoring and additional research– could then be added at later stages in view of developing ever better foresights and risk assessments, and fine-tuning and applying adaptive policies.

A lesson from the 2009 flu pandemic is that nations need to cooperate and better coordinate their vaccination policies: coordination is needed to avoid ‘first come first served’ games played by the major vaccine producers and excessive panic orders being submitted by nation states, and cooperation is needed to start initial vaccine development and production at minimal cost to individual nation states. A sufficient initial order (e.g. 20% coverage) with an option to order more if necessary from a coalition of nations would then lead to the development of vaccines and the start up of production to satisfy the amount initially ordered, and a second amount –informed by monitoring of fatality ratios, infectivity ratios, immunity, et cetera– could be ordered once development is under way and/or production is starting up.

Such a vaccination policy would be an adaptive policy if it would have been pre-specified and a monitoring system would have been put in place. For vaccination policies, adaptive vaccination policies do not seem to exist in most countries yet – as opposed to social distancing policies, which are already adaptive in several countries in the sense that pre-specified sets of measures are activated semi-automatically when specific thresholds are crossed.

Simulations with the flu model presented in this paper show that rich countries should seriously consider the rest of the world when designing and implementing new flu policies – not only for ethical reasons, but also for their own sake – since it takes a systems perspective with feedback between regions.

The A(H1N1)v flu virus is a good example for –but should not be taken as a blueprint for– future outbreaks: new epidemics/pandemics are unpredictable – they could range from very mild to extremely catastrophic. When confronted with the outbreak of a new type of virus, it makes sense to quickly perform model-assisted exploratory analyses and test whether adaptive policies are robust given all remaining uncertainties. The same exploratory models could later be refined and fed with newly available data in order to reduce the uncertain space and apply the adaptive policies given the best knowledge and most recent information available..

5.2 Conclusions regarding ESD and ESDMA

The high-level fast-to-build model described in this paper was developed soon after the first signs of an outbreak of a new flu variant in Mexico, in order to explore possible dynamics of a new flu variant. It is exploratory in terms of its goals, size, scope, and level of aggregation/simplification.

The ESD approach is an interesting quick and dirty assessment approach of plausible dynamics and policy effectiveness. SD and ESD do not differ in terms of the techniques used, but they do somewhat differ in terms of the purposes and issues for which they could be used and the way insights are sought and simulation results are interpreted.

The ESD analyses are nevertheless not deep and broad enough from a ‘deep uncertainty perspective’, because (i) the uncertainty ranges considered are too narrow and conservative, (ii) the analysis of the dynamic behaviours is too shallow, and (iii) robustness of policies is not tested over the entire uncertainty space.

From a methodological point of view it can be concluded that ESD is useful for quickly and somewhat intuitively exploring uncertain dynamically complex issues, e.g. as part of a rapid threat assessment. It may for example be used for generating distinct plausible dynamic behaviours that

may support the reflection process – as ‘tools for thinking’. It may also be used to reason about policies. But in isolation, it may also be criticised for being too intuitive, narrow and unsystematic to base real world policies on.

Later, this simple ESD model was used for ESDMA which allows for a more systematic and broader exploration of uncertainties and policy robustness. From the ESDMA, it can be concluded that the *infected fraction* does not reach the feared third of the population in most runs. However, catastrophic flu pandemics are/were most certainly plausible.

From a methodological point of view it can be concluded that ESDMA is useful for systematically exploring *plausible* behaviours of uncertain dynamically complex issues as well as the efficiency and effectiveness/robustness of potential solutions under deep uncertainty. It may for example be used to systematically scan and analyse all possible behaviours that can be generated by taking parametric and structural uncertainties seriously into account. However, ESDMA takes much more time than ESD. Hence, the type of analysis may also depend on the time available for analysis, or in other words, the urgency of starting to take action (see (Pruyt 2010c)). ESDMA may from that point of view be more appropriate for in-depth risk assessments when time is available to do so.

5.3 Future Work

The exploratory use of SD models turns these models into extremely powerful tools for generating insights and supporting decision making in case of complex issues under deep uncertainty. We will further pursue this line of research. In the near future, we will work on cases of increasing complexity in order to gradually advance our work on ESDMA. In this paper, we explored –on top of uncertainty ranges and different formulations/functions– (slightly) different submodels. Simple analysis and visualisation techniques were good enough here because the behaviours generated with the different variants of the exploratory flu model were sufficiently distinctive. In future research, we will simultaneously explore broad ranges, functions, boundaries –in other words (very) different model formulations– as well as different world views with very different models and preferences. Each of these steps will lead –at first sight and without good analyses and visualisation techniques– to more and more chaotic behaviours. We will therefore also work on developing ways to deal with problems of combinatorial complexity, and develop better and more user-friendly methods, techniques, and tools, starting from existing analysis methods, data mining techniques, mathematical and control theory techniques, formal modelling methods, and multi-dimensional visualisation techniques.

References

- Agusdinata, D. (2008). *Exploratory Modeling and Analysis. A Promising Method to Deal with Deep Uncertainty*. Phd dissertation, Delft University of Technology, Delft. <http://www.nextgenerationinfrastructures.eu/index.php?pageID=17&itemID=433706>. 4
- Bankes, S. (1993, May-June). Exploratory modeling for policy analysis. *Operations Research* 41(3), 435–449. 3
- Bankes, S., W. Walker, and J. Kwakkel (2010). *Encyclopedia of OR* (upcoming 3th edition ed.), Chapter Exploratory Modeling and Analysis. Boston: Kluwer Academic Publishers. [forthcoming]. 3
- Bryant, B. and R. Lempert (2009). Thinking inside the box: A participatory, computer-assisted approach to scenario discovery. *Technological Forecasting & Social Change* 77(2010), 34–49. 4
- Centers for Disease Control and Prevention (2009, April 24). Swine influenza a (h1n1) infection in two children – southern california, march–april 2009. *Morbidity and Mortality Weekly*

- Report 58(15), 400–402. <http://www.cdc.gov/mmwr/preview/mmwrhtml/mm5815a5.htm> (posted on April 21 as an MMWR Early Release on the MMWR website). 1
- ECDC (2009a, 12 June). ECDC interim risk assessment. www.ecdc.europa.eu. 2
- ECDC (2009b, 8 May). ECDC risk assessment. www.ecdc.europa.eu. 2
- ECDC (2009c, 20 May). ECDC risk assessment. www.ecdc.europa.eu. 2
- ECDC (2009d, 20 July). ECDC risk assessment. www.ecdc.europa.eu. 2
- ECDC (2009e, 21 August). ECDC risk assessment. www.ecdc.europa.eu. 2
- ECDC (2009f, 30 April). ECDC threat assessment. www.ecdc.europa.eu. 2
- ECDC (2009g, 24 April). ECDC threat assessment – update. www.ecdc.europa.eu. 2
- Ford, A. (1999). *Modeling the environment: an introduction to system dynamics models of environmental systems*. Washington (D.C.): Island Press. 29
- Forrester, J. (1968). *Principles of Systems*. Cambridge, MA: Wright-Allen Press, Inc. 29
- Kwakkel, J., W. Walker, and V. Marchau (2010, 2010). From predictive modeling to exploratory modeling: How to use non-predictive models for decisionmaking under deep uncertainty. In *Proceedings of the 25th Mini-EURO Conference on Uncertainty and Robustness in Planning and Decision Making (URPDM2010), 15-17 April*, Portugal. University of Coimbra. 3
- Lempert, R., D. Groves, S. Popper, and S. Bankes (2006). A general, analytic method for generating robust strategies and narrative scenarios. *Management Science* 52(4), 514–528. 3
- Lempert, R., S. Popper, and S. Bankes (2003). Shaping the next one hundred years: New methods for quantitative, long-term policy analysis. RAND report MR-1626, The RAND Pardee Center, Santa Monica, CA. http://www.rand.org/pubs/monograph_reports/2007/MR1626.pdf. 2, 3
- Meadows, D. and J. Robinson (1985). *The Electronic Oracle. Computer Models and Social Decisions*. Chichester: John Wiley & Sons. 29
- Pruyt, E. (2007). Dealing with Uncertainties? Combining System Dynamics with Multiple Criteria Decision Analysis or with Exploratory Modelling. In *Proceedings of the 25th International Conference of the System Dynamics Society*, Boston, MA, USA. System Dynamics Society. www.systemdynamics.org/conferences/2007/proceed/papers/PRUYT386.pdf. 4
- Pruyt, E. (2009, July). Cholera in Zimbabwe. In *Proceedings of the 27th International Conference of the System Dynamics Society*, Albuquerque, USA. System Dynamics Society. <http://www.systemdynamics.org/conferences/2009/proceed/papers/P1357.pdf>. 6
- Pruyt, E. (2010a, July). Making System Dynamics Cool II: New hot testing and teaching cases of increasing complexity. In *Proceedings of the 28th International Conference of the System Dynamics Society*, Seoul, Korea. System Dynamics Society. <http://www.systemdynamics.org/cgi-bin/sdsweb?P1026+0>. 4
- Pruyt, E. (2010b). Scarcity of minerals and metals: A generic exploratory system dynamics model. In *Proceedings of the 28th International Conference of the System Dynamics Society*, Seoul, Korea. System Dynamics Society. <http://www.systemdynamics.org/cgi-bin/sdsweb?P1268+0>. 4
- Pruyt, E. (2010c, July). Using Small Models for Big Issues: Exploratory System Dynamics for Insightful Crisis Management. In *Proceedings of the 28th International Conference of the System Dynamics Society*, Seoul, Korea. System Dynamics Society. <http://www.systemdynamics.org/cgi-bin/sdsweb?P1266+0>. 3, 4, 26
- Pruyt, E. and C. Hamarat (2010). The concerted run on the DSB Bank: An Exploratory System Dynamics Approach. In *Proceedings of the 28th International Conference of the System Dynamics Society*, Seoul, Korea. System Dynamics Society. <http://www.systemdynamics.org/cgi-bin/sdsweb?P1027+0>. 4

Sterman, J. (2000). *Business dynamics: systems thinking and modeling for a complex world*. Irwin/McGraw-Hill: Boston. 29

Sterman, J. (2002). All models are wrong: reflections on becoming a systems scientist. *System Dynamics Review* 18(4), 501–531. 2

Trifonov, V., H. Khiabani, B. Greenbaum, and R. Rabadan (2009, April 30). The origin of the recent swine influenza a(h1n1) virus infecting humans. *Euro Surveill* 14(17). <http://www.eurosurveillance.org/ViewArticle.aspx?ArticleId=19193>. 1

A System Dynamics and Diagrammatic Conventions

SD models consist of specific structural elements causally linked into feedback loops. These models only contain *direct causal* relations (e.g. the blue links in Figure 13). For SD to be of any use, it is required that possible causal links can be perceived or hypothesised.

A *feedback loop* consists of two or more causal influences between elements that are connected in such a way that if one follows the causality starting at any element in the loop, one eventually returns to the first element. In other words, the variable feeds back –after some time– to itself, which makes that its behaviour is (partly) shaped by its own past behaviour. Feedback loops are either positive or negative. Positive feedback loops (an even number of negative links) in isolation generate exponential growth or decay. Negative feedback loops (an odd number of negative links) in isolation generate balancing or goal-seeking behaviour and can be used for automatic control/balancing. Feedback loops give rise to nonlinear behaviour, even if all constitutive causal relationships are linear. Feedback loops almost never exist in isolation: several feedback loops are often strongly connected, and their respective strengths change over time. The feedback concept is a fundamentally important characteristic of SD.




Figure 13: A detailed CLD

Causal Loop Diagrams (CLD) are often used to map feedback loop structures. Following symbols are used in a CLD: $\rightarrow$ represents a positive causal influence; $\rightarrow$ represents a negative causal influence; $\rightarrow$ and $\rightarrow$ represent a positive and a negative causal influence with a delay; $\ominus$ and $\ominus$ represent negative feedback loops; and $\oplus$ and $\oplus$ represent positive feedback loops. Figure 13 displays a detailed CLD of the simulation model displayed in Figure 14. The feedback loop structure is often easier to communicate by means of strongly simplified/aggregated CLDs.

SD models also contain *stock-flow structures*. Figure 14 shows a basic Stock-Flow Diagram (SFD), which graphically represents a SD simulation model, and more specifically, its stock variables ($\square$), flow variables ($\Downarrow$), auxiliary variables ($\circ$ or no symbol), and constants ($\diamond$ or no symbol) and other direct causal influences between variables (the blue arrows). A stock variable –also called a level or a state variable– could be seen metaphorically as a ‘bath tube’ or ‘reservoir’. During a simulation, stock variables can only be changed by flow variables (also known as rates).

Every feedback loop contains at least one stock variable or memory (in order to avoid simulation problems caused by simultaneous equations).




Figure 14: A SFD representing a SD simulation model

Positive inflows increase the contents of stock variables, and positive outflows decrease their contents: ingoing flows are, metaphorically speaking, taps or valves, and outgoing flows drains. Flow variables regulate the states of stock variables. Hence, flow variables are the variables that need to be targeted by policies to improve the problematic condition/state of the more inert stock variables. Mathematically speaking, the stock variable is the integral of the difference between the incoming flows and the outgoing flows over the time interval considered, plus the amount in the stock at the beginning of the period.

Time *delays* are also important elements of SD models. They are included in causal loop diagrams by means of slashed arrows ($\rightarrow$), and in stock/flow diagrams by different delay-type functions and slashed arrows ($\rightarrow$). *Nonlinear functions* may also be important in SD models. They are often included in computer models by means of (non-linear) table functions, also called lookup functions or graph functions.




Figure 15: Behaviour of a SD simulation model

The simulation over time of simulation models of these structural elements gives what system dynamicists are really interested in: the overall modes of behaviour (see Figure 15). System Dynamics is not to be used for exact point prediction or path prediction (Meadows and Robinson 1985, p34). One of the basic assumptions of SD is that the structure of a system (and model) drives its behaviour, and hence, that structural policies are needed to effectively and robustly change (possible) undesirable behaviours. System dynamicists are therefore biased towards structural solutions by trying to structurally change (information) links, feedback loops, underlying stock-flow structures, etc.

For introductions into the SD field, see among else (Sterman 2000) (Ford 1999) (Forrester 1968).

B The Automated ESD Model




Figure 16: Stock-flow diagram of the ESD two-region flu model with the core structure for region 1 in bold (basic variables are displayed in black connected by dark blue arrows, initial values and parameters in orange, switches to turn structures of/off in light blue and yellow, and policy structures in green)

C ESDMA Results




Figure C.1: Infected fraction (region 1) for the entire ensemble of plausible scenarios for the FFD (5x5 graphs with $5*5*5*5*3 = 1875$ trajectories per graph)




Figure C.2: Detail of the infected fraction (region 1) for infection rates of 0.10 (both regions 1 and 2)




Figure C.3: 3D plot and projections for LH sampling varying 17 uncertain parameters (20000 runs)




Figure C.4: Policy 1 (static) Between square brackets the fraction of runs with a total fatality rate of more than 0.1% (600.000 out of 600.000.000)




Figure C.5: Policy 2 – static contact rate reduction [Between square brackets the fraction of runs with a total fatality rate of more than 0.1%, 1% and 2.5%]




Figure C.6: Policy 3 (static). [Between square brackets the fraction of runs with a total population fatality rate of more than 0.1% (600.000 out of 600.000.000)]