Witness Name: Tim Savage

Statement No.: 1

Exhibits: 0

Dated: 07 August 2023

UK COVID-19 INQUIRY

WITNESS STATEMENT OF TIM SAVAGE

I, Tim Savage, will say as follows: -

1. I am the Founder and CEO of Armakuni (Ak). Ak is a technology consultancy which focuses on solving problems using cutting-edge software technologies and methods, in a collaborative approach with our clients. We work with both the public and private sectors.

2. At the end of the work under discussion, all data and artefacts including model code, output documents, and dashboards were fully handed over to Welsh Government ownership, sitting on cloud servers which they own; we no longer have copies of these. The details of this statement are therefore reconstructed from the best recollection of the project team, and from the contractual/legal and email communications that we do still have accessible.

3. In April 2020, we were commissioned to urgently support the Welsh Government (WG) with improving their Covid situational awareness and decision-making support tools. Our team and skillset were familiar to them due to work we had been doing for NHS Wales Informatics Services, which is a partner agency of WG. WG approached us to see if we could help with some of the challenges they were facing in their response to the Covid epidemic. The work was commissioned through the G-Cloud 11 framework. The two major pieces of work were:

4. Technical development and delivery of a data dashboard showing a cross-agency view for Wales and its regions, in as near real-time as possible, of key

Covid-related issues, including Covid transmission, cases, and deaths, and NHS capacity. The initial focus of this work was to create a single source of information in order to give decision makers a common, accessible and accurate view; the overall brief was that "a policy maker should be able to understand everything they need to make key decisions within 2 minutes". Over time, this work evolved to have more emphasis on early warning indicators and 'circuit breaker' status, to enable a clearer consensus view on risk levels. WG officials worked with their agency counterparts to decide what data should be included and then shared those requirements with us. Our responsibility was to work out how to acquire, combine, and present this data effectively.

5. Implementation of a Wales-specific instance of a pre-existing model, OpenABM-Covid19, which aimed to simulate the expected impact of a variety of NPIs (non-pharmaceutical interventions) under consideration by WG on Covid infections, hospitalisations, and deaths. WG selected the model to be used; our role was to work out the details of implementation, acquire and use the right 200+ data inputs for the model, communicate the model outputs in an understandable way and with appropriate confidence levels, and to test and cross-check our work. The bulk of this work happened in June-August 2020, with some follow-on into October 2020.

6. We worked closely with civil servants during the development of these two products, sharing draft versions and adapting the content and presentation as WG's needs and the Covid situation evolved. The main teams we worked with day-to-day were the Technical Advisory Cell (TAC) and Knowledge and Analytics Service (KAS). At project close, we provided computer code and documentation for the modelling to these officers and to their partners at Swansea University, with the intention that Swansea could take development forward in-house if needed in future.

7. As is usual in our work, throughout the project we also provided support and skills development for the WG team on how to rapidly deliver technical products using an agile approach.

8. In terms of communication of data and statistics, WG decided what data and statistics should be used and in relation to which questions. We were then

responsible for designing appropriate visualisations and interfaces which clearly communicated that information, on both the dashboard and the modelling work.

9. Modelling in more detail

10. Representatives of WG were attending SPI-M meetings, but were finding that there was insufficient discussion or data relating to Wales specifically for their needs. Therefore, there was a need for modelling which focussed on Wales and the possible impacts of policy interventions there, especially where transmission of Covid or policies might differ from England.

11. We were asked to apply an existing model which was already in use in England for the same purpose, OpenABM-Covid19 ("the Oxford Model"), to a Welsh context, We were sense-checking, calibrating, providing appropriate Wales-specific input data, and developing ways to run the model rapidly to explore and share the outcomes of multiple policy scenarios.

12. The original core model was developed by the University of Oxford. The model's structure and assumptions are well described on its web page.

13. During the project, we worked with a number of other individuals and organisations who were involved in parallel initiatives, to get up to speed as fast as possible and to share approaches to model calibration and implementation, namely:

14. NHS England/NHSX - getting information from Jonny Pearson on NHSE/X's use of the same modelling approach

15. Oxford Big Data Institute - a team led by Christophe Fraser, who had developed the original model, to understand its basic history and rationale. We also later interacted with the project team via email several times to clarify details such as data definitions, and to find out about the timelines for and contents of intended updates

16. Faculty.ai - we spoke on several occasions to Scott Stevenson and his technical colleagues at Faculty. Their team had been tasked with refining the OpenABM-Covid19 model for an English context, and we met to compare

mathematical methods and approaches for efficiently running, calibrating, and testing the model

17. Near the project's completion in September 2020, we met with representatives of Swansea University, led by Professor Biagio Lucini. At that time, WG were interested in the potential of Swansea continuing to provide modelling capability and development longer term, so wanted us to provide a technical handover to them

18. We understand that WG were using several different approaches alongside OpenABM to triangulate their decisions, though we were not informed of the detail of the other approaches.

19. OpenABM is an agent-based model, a simulation of simplified human behaviour, which observes directly what happens when a few million 'agents' with given characteristics are allowed to mix and potentially spread infection under different scenarios. This type of model is useful in that it allows to directly inspect what might emerge from widespread behaviour changes, for example changes in the number of contacts by children if schools were closed; this can be much harder to work out in non-agent-based modelling structures, as it's so hard to understand how the subtleties of, for example, the ages of the various actors may have an impact on end results. Another strength of this approach is that once initially set up, it allows large numbers of empirical test runs of different policies and combinations of policies.

20. Like all models, this approach also has limitations. The most obvious in this setting are that:

21. The model simplifies human characteristics and behaviour down to a few dimensions. In this case, the main dimensions were household size, individual's age, and average levels of social contact. Though we were not the designers, we would infer that age is a foundational input of the model both because it was known at that time to be a highly significant variable in the risk of infection, severe illness, and death, and because some of the interventions being considered had unique impacts on some age groups, e.g. school or workplace closures.

22. The model also has simplifying assumptions to make the results computable in a reasonable amount of time - for example, there was no spatial aspect to the modelling, and the model did not account for people entering or exiting the country.

23. Many fundamentals of Covid transmission were unknown or uncertain at that stage, for example the proportion of infected individuals who were asymptomatic; the modelling was based on the best recent data available within the UK at the time, but that data was rapidly evolving.

24. Once the model was up and running and fully checked, WG officers shared a number of potential policy scenarios of interest with us, and we ran the model to create outputs estimating the outcomes of these policies. The first outputs which simulated future policy outcomes were shared with WG in late July 2020, and to our understanding were used in meetings and discussions with policy decision makers immediately thereafter.

25. From initial use onwards, there were not significant changes in the model structure, though the BDI team did release a version 0.3 in July 2020 with some minor improvements. The main aspects of our work were in improvements to inputs and outputs, namely:

26. Updated input assumptions, for example updated data from the London School of Hygiene and Tropical Medicine's CoMix study, which gave us data on average numbers of social contacts per day during lockdown conditions, and regularly updated data on daily hospitalisation figures. These changes were made to improve the expected accuracy of the model and its simulation of policy alternatives.

27. Improved output formats and approach, developing a systematic process to record the desired policy tests, rapidly run the experiments, and create a standard package of results tables and graphs which we could email over immediately. These changes were made to make it faster and easier for us to respond to new simulation requests in a repeatable way, and in a format which WG officers could quickly absorb and communicate to decision-makers.

28. In general, there was sufficient data available to develop the model to a usable level of quality - as expected, given that England had already gone forward with this approach. Basic demographic data was very easy to obtain from Stats Wales and was of high quality; live data on Covid hospitalisations and deaths was a less smooth process but was rapidly compiled and provided to us by WG officers. Other base parameters, for example the % of asymptomatic infections, were much harder to know or estimate, but we had access to the best estimates available from SPI-M which we used as our working assumptions. These limitations would have had some impact on the reliability of the model, but at the time there were gaps in data for any modelling approach which was available.

29. During the relevant period, we were in regular communication with WG TAC and KAS officers, and we often requested that they source data which we expected to have significant impact on the model outputs. For example, they acquired an updated set of CoMix data on our behalf as the accuracy of this data would make a significant difference to the modelling. Much of our discussion with them was about available sources of input data, its definitions and quality, and to what extent this would affect the ability to build useful simulations.

30. The mathematical modelling principally took account of the impact of Covid-19 on groups split by age, as described above; this was because age was structurally very impactful in Covid transmissions and policy interventions. It did not differentiate between impacts on at-risk or vulnerable groups in any other way.

31. Likewise, in the dashboard work we undertook, we were not asked to and did not take specific account of at-risk and vulnerable groups; we were focussed on Wales-wide situational awareness and impacts, and were not asked to split this out by any specific group of people.

32. Dashboard production in more detail

33. As well as the modelling work, Ak were commissioned to develop a dashboard (not a platform) bringing together surveillance data from multiple partners. This was in order to enable decision-makers to have an easily accessed shared view of key Covid indicators.

34. WG officers decided what data and themes to include in the dashboard work. The selection of data themes was driven by the need to understand the rate of change of severe illness or death across each of the health boards in Wales, and how close the situation was to reaching health services capacity, for example available ICU beds. This was to enable discussion and decision-making about when or whether it would be necessary to enact more dramatic 'circuit breaker' policies. We understand that WG conducted workshops with multiple agencies to work out what data should be included, based on all agencies' needs.

35. Data collection in order to build the dashboard was relatively challenging, which we believe is part of why we were brought in. The data came in a large number of different forms, including PDFs, spreadsheets, from APIs, and scraped from websites. Source data also had frequent changes in format, making it difficult to automate, and definitions of data and who was the owner was often very unclear.

36. This meant that having a reliable daily dashboard update was difficult, and that there was a risk that some data was misleading, for example if it had undeclared quality issues.

37. In our view, the main obstacles or missed opportunities were that the variety of formats and locations of source data, and the ambiguities around ownership and SLAs for data provision, meant that data was harder to put together and WG could rely on it less confidently than they otherwise might. This also meant that public servants were often occupied with manual tasks such as calling around hospitals for case numbers, or chasing down missing datasets, when they might otherwise have been focussed on higher-value efforts.

38. A lack of clarity on source data formats, definitions, quality, and ownership were certainly factors which slowed down the delivery of a shared data picture in mid-2020, although we were not involved earlier than May 2020 so can't comment on that time period in particular. Human factors stemming from a lack of formal clarity were the major issue - because the various agencies were not in the habit of working cross-functionally, and didn't have routine working relationships with defined responsibilities to one another, it was difficult to resolve data sourcing issues when they did arise, or to enact improvements where those would be sensible.

39. Data governance was less of an issue in this case; no personally-identifiable data was needed for either the modelling or the dashboard work, and most of the data being sourced was not highly sensitive so governance questions were usually straightforward to resolve.

40. Data linkage was not an issue, as we did not need to link across different datasets for any of this work.

41. The lack of a data platform and of data engineering capacity had been a limiting factor for WG, which is why they engaged Ak. If WG had more internal technical capacity, or if the datasets had been easier to work with, they would have been more likely to be able to build this without our assistance.

42. There was additional complexity because of the differing data systems across organisations. Authentication and authorisation was different for each organisation, including needing access to particular physical hardware or VPNs, and there were a number of different visualisation softwares in place with different access licences, e.g. PowerBI and Tableau - driving the need for Ak to build a single solution which was accessible across organisations.

43. While we would not make a case for a single cross-Wales data platform, as that raises a number of other challenges, it would certainly make sense in future to have better common standards for authentication, authorisation, and data supply formats and SLAs.

44. As is common practice in agile software development, we undertook regular internal 'retrospectives' during the project duration and at completion, aimed at reflecting on what went well, what did not, and what we would change in the next period of work.

45. While most of these reflected on making our own technical and organisational practices more effective, it was a major theme of our time with WG that interoperability and workflows between teams and agencies were a constant issue, affecting the ability to do work repeatedly and efficiently, at a high level of quality and at a fast pace. If there was one lesson to focus on, we would consider that improving this situation would pay dividends to WG and its partner agencies, both in day to day working and even more in any future crisis.

46. The only UK Government organisation that we liaised directly with was NHS England, in order to understand the genesis and usage of the OpenABM model. In addition, we liaised in more technical detail with NHSE's suppliers, Faculty.ai, in order to share and mutually improve our approaches to deploying copies of the same model.

47. Beyond that, we were using data supplied from UK Government organisations via WG in order to populate our model parameters and the dashboard, for example the latest consensus R-values from SPI-M, but we were not liaising or sharing information directly. These were often shared as slides or PDFs rather than data feeds, but given that it was relatively small amounts of data, this was not in general a major problem.

48. It is possible that with better England/Wales co-ordination, some re-work might have been avoided - for example, we understand that the Cabinet Office had a similar dashboard to the one we built, and we know that NHSE had a parallel model to the one we built - but it is possible that combining these efforts more closely might have created unacceptable reductions in flexibility for the respective national governments.

49. We were not asked to use data about or to analyse the issue of inbound international travel.

Statement of Truth

I believe that the facts stated in this witness statement are true. I understand that proceedings may be brought against anyone who makes, or causes to be made, a false statement in a document verified by a statement of truth without an honest belief of its truth.

Signed: Personal Data

Dated: 07 August 2023