University of Bath

Alternative formats
If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

Operational research for the safe and effective design of COVID-19 mass vaccination centres

Authors:

Richard M Wood [a, c]

Ben J Murch [a]

Simon J Moss [a]

Joshua MB Tyler [a]

Alexander L Thompson [b]

Christos Vasilakis [c]

Affiliations:

[a] Bristol, North Somerset and South Gloucestershire Clinical Commissioning Group, UK National Health Service, Bristol, UK.

[b] Institute for Risk and Disaster Reduction, University College London, London, UK.

[c] Centre for Healthcare Improvement and Innovation, School of Management, University of Bath, Bath, UK.

Correspondence: Dr Richard M Wood.

NHS Bristol, North Somerset and South Gloucestershire Clinical Commissioning Group, South Plaza, Marlborough St, Bristol, BS1 3NX, UK.

richard.wood16@nhs.net.

0000-0002-3476-395X.

Operational research for the safe and effective design of COVID-19 mass vaccination centres

With news, in late 2020, that vaccination against COVID-19 may be up to 95% effective, we have entered a new chapter in our fight against the disease [1]. Restrictions on movement and social contacts can recede as vaccine-acquired immunity reduces susceptibility to infection and (possibly) transmission. A key determinant of this is the speed at which the population can be vaccinated. To facilitate rapid dissemination, many countries are considering mass vaccination centres [2]. Ideally located in large spaces – conference venues or sporting arenas – these sites will immunise hundreds or possibly thousands of individuals each day.

Crucial to their success is the safe and effective planning of demand and capacity. If more people are booked than can be seen then large and unmanageable queues will form; compromising social distancing and reducing the likelihood that people will return for their second and final inoculation. If, on the other hand, demand is too far exceeded by capacity then resources are not fully utilised; wasting vaccine and unnecessarily delaying the regression of economically-punitive social restrictions. The question is therefore, how can we safely and sustainably maximise the throughput of these sites?

To this end, Operational Research (OR) can be a valuable asset. Containing a range of practically focused and mainly quantitative methods, OR has a track record in addressing questions of this very nature. While OR techniques have a history within the immunology field, e.g. in strategically optimising the extent to which influenza vaccine should be stockpiled or reactively purchased [3], the more operational question posed here is perhaps better paired to experiences from the healthcare setting – where modelling patient throughput along some kind of 'pathway' is commonplace [4]. In applying such models to the mass vaccination 'pathway' considered here, we demonstrate an example of the contribution that OR can make in this important next stage of our fight against COVID-19.

Live exercise

On 2 December 2020, a live exercise was conducted at one of the sites planned to operate as a mass vaccination centre in the UK. The purpose of Exercise Panacea was to provide a safe learning environment within which to explore processes for administering COVID-19 vaccine to what would be over 1000 people each day. Such exercises are core to health emergency preparedness, supporting the identification of gaps in plans and processes [5]. The exercise involved 'flowing' a number of 'players' through the site, with services provided by members of the 'cast'. Specifically, 70 players were provided with a unique script for each attendance to ensure a range of presentations were considered, e.g. the representation of elderly people or those with hearing or mobility limitations, as well as those with adverse reactions to vaccination.

Exercise Panacea took place at Ashton Gate football/rugby stadium in Bristol (UK), where a large rectangular interior hall normally used for spectator catering and entertainment had, in the days before the exercise, been converted to a space in which the four activities necessary within the mass vaccination process could be performed (Figure 1). While hitherto unconfirmed, it was a consideration at the time of the exercise that, when live, there would be 1560 arrivals per 12-hour operating period facilitated by six registration assistants, 12 clinical assessors, six immunisers, and 64

1   seats for post-vaccination observation, and maximum safe waiting space for six vaccinees before

2   clinical assessment and 15 before vaccination.

3

4   The overarching vaccination process design had been informed by centrally-produced planning

5   guidance (unpublished) suggesting that immunisers work in teams of two and according to fixed

6   staffing ratios to clinical assessment. Recommendations were that each two-immuniser 'pod' could

7   support a throughput of 520 vaccinations per 12-hour operating period (thus 1560 for six immunisers).

8   Ultimately, the number of pods was limited to six due to spatial constraints of the Ashton Gate site.

9   This also restricted the waiting space within the queues for clinical assessment and vaccination.

10

11   While some of these operational parameters had been informed by an earlier live exercise (Exercise

12   Asclepius, the only other live exercise of its kind before Exercise Panacea), there remained

13   uncertainties given the novelty of the operation and the intricacies of each vaccination centre

14   (specifically with regard to the physical layout of the site and the type and training of staff). Indeed, a

15   key objective of Exercise Panacea was to test performance under such a configuration. Yet a robust

16   appraisal was not fully possible, not least since only a third of the envisaged operating capacity was

17   used during the exercise. In these situations, computer modelling can be a valuable asset in addressing

18   any such gaps in understanding and considering 'what if' scenarios not possible to examine in real life

19   [6].

20

21

22   ***Computer simulation modelling***

23

24   Analysis was performed using a versatile open source simulation tool that had been previously been

25   developed by the authors for modelling patient pathways [7,8]. The tool employs a discrete event

26   simulation method which is well-established in healthcare modelling [9]. This works by simulating

27   the real-life events of vaccinees arriving at the centre, queueing (as necessary), and starting and

28   finishing the various activities along the vaccination pathway (Figure 1). These events are generated

1   according to a given arrival rate and the capacity and service time distributions of each activity (i.e.

2   the model inputs). Simulation outputs, calculated by running multiple (1500) replications of the

3   model, relate to the activity-level numbers of vaccinees in service and in queue over time.

4

5   With the aforementioned 'baseline' arrival rate and capacity allocations, what remained was to

6   estimate the durations of time vaccinees would spend at each of the four activities (Figure 1). This

7   was achieved by fitting the appropriate statistical distributions to data collected from the exercise

8   (using maximum likelihood estimation with selection through Akaike Information Criterion [10]).

9   The distribution of registration time was found to be fairly symmetric, and best approximated by a

10 Weibull distribution with a mean and median of 62 seconds. Both clinical assessment and vaccination

11 times were right-skewed and best approximated by a lognormal distribution with a mean and median

12 of 219 and 200 seconds for the former and 187 and 171 for the latter. Observation time was fixed at

13 15 minutes as per the latest guidance. (For more information on the distribution fitting process refer to

14 the Supplementary Material.)

15

16 Simulation results indicate that the baseline allocation is unviable, with a bottleneck forming at the

17 vaccination activity as characterised by a very high number in service (c.f. capacity of six) and an

18 ever-increasing queue (Table 1, Baseline). This finding is, in fact, evident without modelling – an

19 hourly arrival rate of 130 (i.e.1560 over 12 hours) simply cannot be sustained by a pathway

20 containing an activity whose maximum hourly throughput is only 116 (i.e. six immunisers with 187

21 second estimated mean service duration).

22

23 The solution is either to increase capacity or reduce arrivals. With an operational constraint limiting

24 the number of immunisers to no more than six, the arrival rate could be lowered to the level of

25 maximum throughput. While, at first thought, this may seem a reasonable mitigation, it does not

26 appreciate the impact of variability in service duration, which can contribute to the formation of large

27 queues. Although these are smaller than under the Baseline scenario, they still lead to breaches in the

28 15-space waiting area (Table 1, Scenario 1).

1

2   In order to safely accommodate the various peaks and troughs in service duration, the arrival rate

3   should be sufficiently less than maximum throughput [11]. Lowering the arrival rate by 10% (i.e.

4   from 1386 to 1247 over 12 hours), results in performance within operational limits (Table 1, Scenario

5   2). It would, however, be prudent to increase the waiting space for vaccination (from 15), in order to

6   absorb any potential 'shocks' relating to periods of elevated demand or staff shortages. Given spatial

7   constraints of the site, this can be achieved by shifting the vaccination space into a reduced-capacity

8   observation space (noting that observation capacity can be safely reduced since it is considerably

9   under-utilised – as shown in Table 1, the upper 95% CI for number in service (32.3) is approximately

10 half the allocated capacity (64)).

11

12 Registration and clinical assessment are also under-utilised, implying uneconomic use of available

13 resource. Modelling a one-sixth capacity reduction (i.e. to five and ten workers respectively) is not

14 shown to have an adverse performance impact (Table 1, Scenario 3); with the possible opportunity to

15 safely make further reductions, particularly to registration capacity.

16

17

18 ***Concluding remarks***

19

20 Poor management of demand and capacity can result in suboptimal use of resources and excessive

21 queueing. If available waiting space is breached then safety may be compromised as social distancing

22 cannot be maintained. Modelling and computer simulation can provide useful insights to improve the

23 design and operational management of mass vaccination centres.

24

25 The modelling presented here has directly informed operations at the Ashton Gate site. Following our

26 recommendations, the centre went live on 11 January 2021 with an expanded vaccination queueing

27 area and with 1247 vaccinees booked to each 12-hour operating period (i.e. 416 vaccinees per two-

28 immuniser 'pod'). Site management have reported that, with such an arrival rate, a good balance

appears to have been struck between maximising throughput and ensuring patient safety. As such, daily bookings were based upon the 1247 figure for the first six weeks of operation – a time in which any negative patient experience could have generated poor publicity and impacted upon the high levels of public confidence required to ensure good attendance.

Beyond the analysis contained here, future work should more formally assess the impact of unforeseen 'shocks' to the vaccination process. In addition to capturing variation in arrivals and service durations (as in this study), it would be prudent to consider the resilience of any setup to the range of 'low-frequency, high-impact' stochastic events that could be possible. For instance, staff unavailability or a road traffic accident that causes delays followed by a deluge of arrivals. Modelling could be useful in determining the necessary 'slack' in capacity required to safely absorb such shocks.

Given the aforementioned intricacies of each vaccination centre, a 'one-size-fits-all' blueprint would unlikely be appropriate. Instead, those involved in setting up and managing different sites should consider the use of bespoke modelling to initialise or optimise their operation. The simulation tool used here is freely available to such ends [7,8]. With this software, prospective users can experiment with different arrival rates and capacity configurations. The software also has additional functionality to account for time-dependent arrival rates and capacities (for instance, for use in modelling the previously mentioned shocks).

As well as demand and capacity management, OR can contribute to effective mass vaccination in a number of other ways. These may include workforce scheduling, predicting no-shows and associated airline-style 'overbooking', and optimising the priority order of individuals for vaccination based upon their risk of severe illness (older people) and/or onward transmission (younger people).

References

[1]   Polack FP, Thomas SJ, Kitchin N, Absalon J, Gurtman A, Lockhart S, Perez JL, Pérez Marc G, Moreira ED, Zerbini C, Bailey R. Safety and efficacy of the BNT162b2 mRNA covid-19 vaccine. New England Journal of Medicine. 2020 Dec 10. https://doi.org/10.1056/NEJMoa2034577.

[2]   Majeed A, Molokhia M. Vaccinating the UK against covid-19. BMJ. 2020. 371:m4654. https://doi.org/10.1136/bmj.m4654.

[3]   Grieco L, Panovska-Griffiths J, van Leeuwen E, Grove P, Utley M. Exploring the role of mass immunisation in influenza pandemic preparedness: a modelling study for the UK context. Vaccine. 2020 Jul 14;38(33):5163-70. https://doi.org/10.1016/j.vaccine.2020.06.032.

[4]   Palmer R, Fulop NJ, Utley M. A systematic literature review of operational research methods for modelling patient flow and outcomes within community healthcare and other settings. Health Systems. 2018 Jan 2;7(1):29-50. https://doi.org/10.1057/s41306-017-0024-9.

[5]   Skryabina E, Reedy G, Amlot R, Jaye P, Riley P. What is the value of health emergency preparedness exercises? A scoping review study. International journal of disaster risk reduction. 2017 Mar 1;21:274-83. https://doi.org/10.1016/j.ijdrr.2016.12.010.

[6]   Griffiths JD, Williams JE, Wood RM. Modelling activities at a neurological rehabilitation unit. European Journal of Operational Research. 2013 Apr 16;226(2):301-12. https://doi.org/10.1016/j.ejor.2012.10.037.

[7]   The Health Foundation. Developing a versatile tool for modelling pathway capacity in NHS organisations. https://www.health.org.uk/improvement-projects/developing-a-versatile-tool-for-modelling-pathway-capacity-in-nhs-organisations. Last accessed on 15 December 2020.

[8]   NHS BNSSG CCG Modelling and Analytics. PathSimR model code repository. https://github.com/nhs-bnssg-analytics/PathSimR. Last accessed on 15 December 2020.

[9]   Pitt M, Monks T, Crowe S, Vasilakis C. Systems modelling and simulation in health service design, delivery and decision making. BMJ quality & safety. 2016 Jan 1;25(1):38-45. http://dx.doi.org/10.1136/bmjqs-2015-004430.

[10]   Akaike H. Information theory and an extension of the maximum likelihood principle. InSelected papers of hirotugu akaike 1998 (pp. 199-213). Springer, New York, NY. https://doi.org/10.1007/978-1-4612-1694-0_15.

[11]   Kendall DG. Some problems in the theory of queues. Journal of the Royal Statistical Society: Series B (Methodological). 1951 Jul;13(2):151-73. https://doi.org/10.1111/j.2517-6161.1951.tb00080.x.
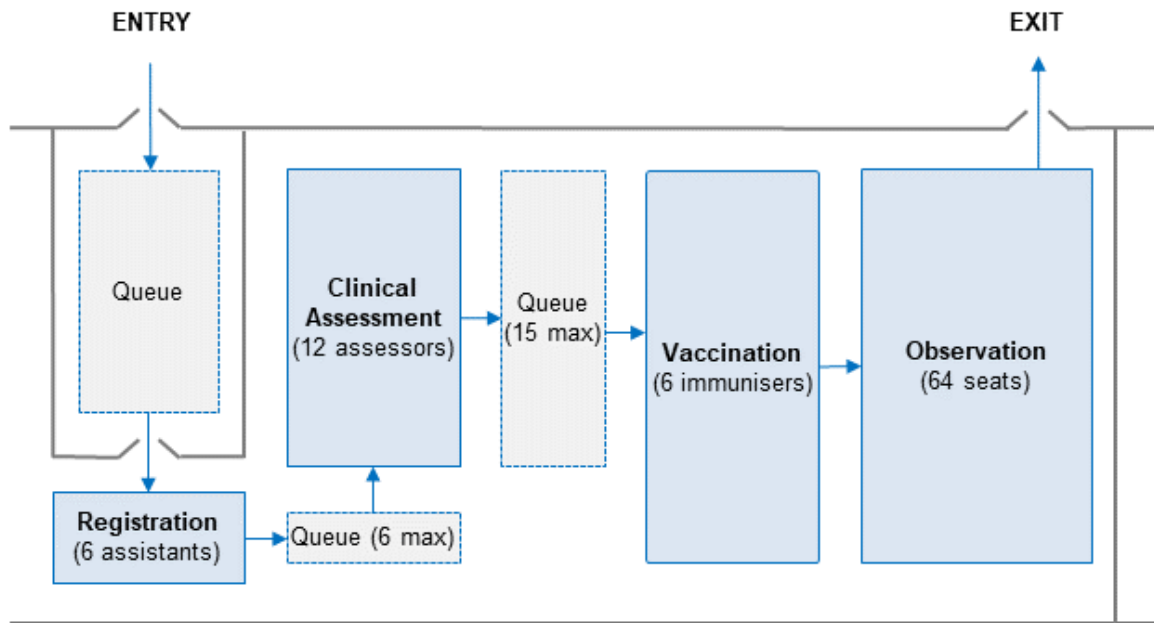
Figure 1. Configuration of the Ashton Gate mass vaccination centre in Bristol, where arriving individuals pass through four activities: registration, clinical assessment, vaccination and observation. If any of these stages are full, then individuals queue in the dedicated waiting areas.

Table 1. Steady-state simulation results for number of vaccinees in service and in queue under the Baseline scenario and hypothetical Scenarios 1 to 3. *Arrivals* is the number of vaccinees arriving at the site per hour and *Capacity* represents the maximum number of vaccinees that can concurrently be served within registration, clinical assessment and vaccination respectively. Note, unless otherwise indicated, steady state was reached within the first hour of the 12-hour operating period.

Scenario	Arrivals	Capacity	Mean number (95% CI) of vaccinees in service				Mean number (95% CI) of vaccinees in queue			
			Registration	Clinical assessment	Vaccination	Observation	Registration	Clinical assessment	Vaccination	Observation
Baseline	130	6-12-6	2.2 (0.0 – 5.6)	7.8 (3.0 – 12.0)	6.0 (6.0 – 6.0)*	29.6 (24.7 – 34.4)	0.0 (0.0 – 0.0)	0.2 (0.0 – 2.2)	108.0 (34.9 – 179.0)*	0.0 (0.0 – 0.0)
Scenario 1	116	6-12-6	2.0 (0.0 – 5.0)	6.9 (2.3 – 12.0)	5.7 (3.0 – 6.0)**	28.1 (20.3 – 33.7)	0.0 (0.0 – 0.0)	0.1 (0.0 – 0.6)	9.1 (0.0 – 36.5)**	0.0 (0.0 – 0.0)
Scenario 2	104	6-12-6	1.8 (0.0 – 4.9)	6.2 (2.0 – 11.6)	5.1 (1.7 – 6.0)	25.4 (16.8 – 32.3)	0.0 (0.0 – 0.0)	0.0 (0.0-0.0)	2.1 (0.0 – 11.6)	0.0 (0.0 – 0.0)
Scenario 3	104	6-10-5	1.8 (0.0 – 4.9)	6.2 (2.0 – 10.0)	5.1 (1.8 – 6.0)	25.4 (16.9 – 32.4)	0.0 (0.0 – 0.0)	0.1 (0.0-1.9)	2.1 (0.0 – 11.1)	0.0 (0.0 – 0.0)

* Values at end of 12-hour operating period. Behaviour did not stabilise during operating period.

** Values from hours 8 to 12 within operating period. Behaviour stabilised at approximately hour 8.

Supplementary Material

Duration of time at registration, clinical assessment and immunisation

1. Background

The main steps of the mass vaccination process from an operational perspective are:

1. Registration at front desk (S1)
2. Clinical assessment (S2)
3. Immunisation (S3)
4. Observation (S4)

Quantifying the duration at each step of the process is useful for effective planning and is an important prerequisite for understanding capacity sufficiency via dynamical modelling. Of crucial importance is the distribution of such durations. Such variability is essential to factor into operational considerations, since larger variability puts greater pressure on queue holding areas and increases waiting time[1].

From the live exercise, various times were recorded for durations at S1 to S3 above. However, these are just samples, and are moreover samples with a relatively low sample size. In estimating the truer distribution of the S1-S3 durations from the underlying population, the appropriate statistical distributions can be fitted to the sample data.

Here, a number of candidate distributions (known to perform well in healthcare settings[2] [3]) are fitted and the most appropriate one is selected based upon AIC[4]. Note that S4 is not considered since this is assumed to be of constant 15 minute duration.

Each section in this report contains results for each step S1-S3, illustrating the quality of the distribution fit to the data and presenting corresponding estimates for mean and median duration, and standard deviation, as derived from the fitted distribution. Full results are in the appendix.

[1] Gupta D. Queueing models for healthcare operations. InHandbook of Healthcare Operations Management 2013 (pp. 19-44). Springer, New York, NY.

[2] Marazzi A, Paccaud F, Ruffieux C, Beguin C. Fitting the distributions of length of stay by parametric models. Medical care. 1998 Jun 1:915-27.

[3] Faddy M, Graves N, Pettitt A. Modeling length of stay in hospital and other right skewed data: comparison of phase-type, gamma and log-normal distributions. Value in Health. 2009 Mar;12(2):309-14.

[4] Akaike H. Information theory and an extension of the maximum likelihood principle. InSelected papers of hirotugu akaike 1998 (pp. 199-213). Springer, New York, NY.

2. Registration at front desk (S1)

Best distribution is Weibull with shape and scale parameters 3.92 and 68.53 respectively.

Figure 1. Illustration of Weibull distribution fitted to sample data (note: units in seconds).

Table 1. Comparison of sample and fitted distribution moments (note: units in seconds).

	Median	Mean	Standard deviation
Sample (n=20)	65	62	19
Distribution	62	62	18

3. Clinical assessment (S2)

Best distribution is lognormal with shape and scale parameters 0.429 and 5.30 respectively.
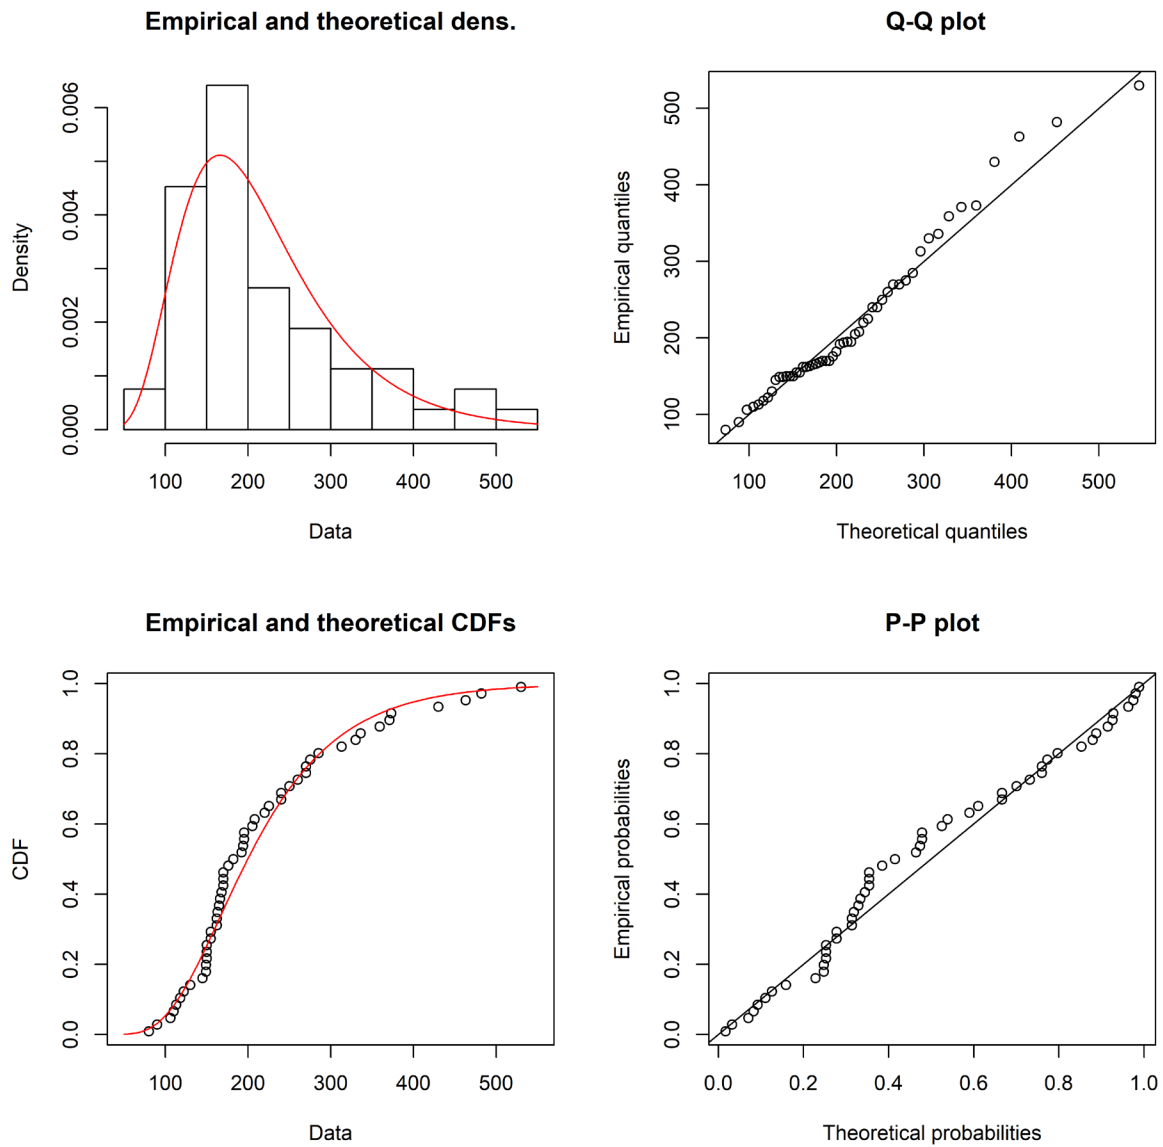

Figure 2. Illustration of lognormal distribution fitted to sample data (note: units in seconds).

Table 2. Comparison of sample and fitted distribution moments (note: units in seconds).

	Median	Mean	Standard deviation
Sample (n=53)	182	220	103
Distribution	200	219	98

4. Immunisation (S3)

Best distribution is lognormal with shape and scale parameters 0.414 and 5.14 respectively.
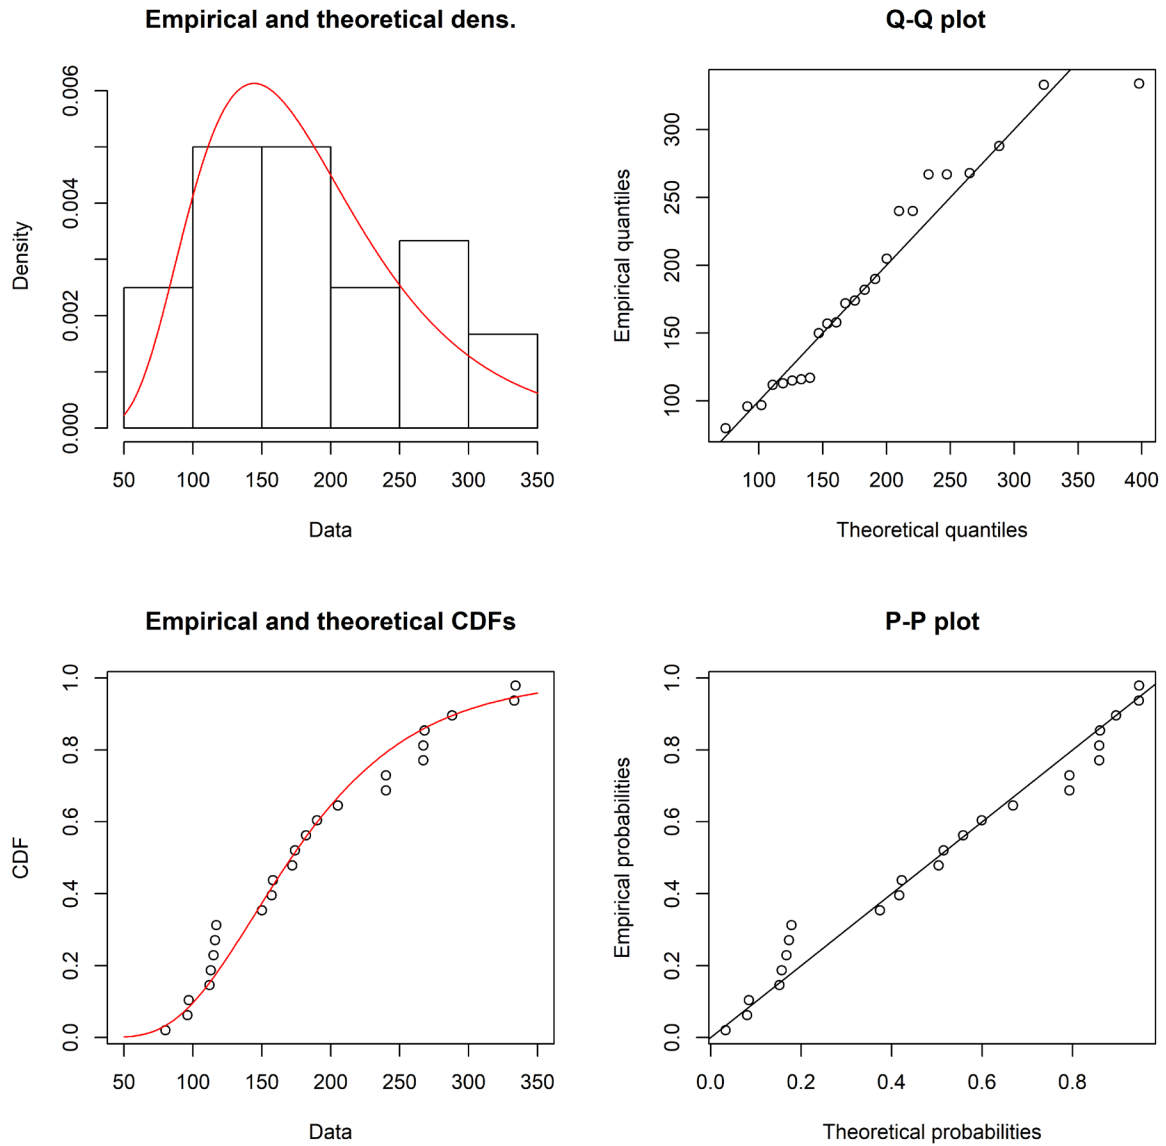

Figure 3. Illustration of lognormal distribution fitted to sample data (note: units in seconds).

**Table 3.** Comparison of sample and fitted distribution moments (note: units in seconds).

	Median	Mean	Standard deviation
Sample (n=24)	173	186	77
Distribution	171	187	81

Appendix A.1 Registration at front desk (S1)

Histogram and theoretical densities

Empirical and theoretical CDFs

Q-Q plot

P-P plot

	expo	normal	tnormal	lognormal	gamma	weibull
Kolmogorov-Smirnov statistic	0.372	0.158	0.158	0.224	0.204	0.148
Cramer-von Mises statistic	0.924	0.067	0.067	0.142	0.113	0.059
Anderson-Darling statistic	4.553	0.347	0.347	0.743	0.580	0.318
Akaike's Information Criterion	207	177	177	181	179	176
Bayesian Information Criterion	208	179	179	183	181	178

eCDF comparison, with bootstrapped CIs

Cumulative probability

Duration in service point (seconds)

Appendix A.2 Clinical assessment (S2)

Histogram and theoretical densities

Empirical and theoretical CDFs

Q-Q plot

P-P plot
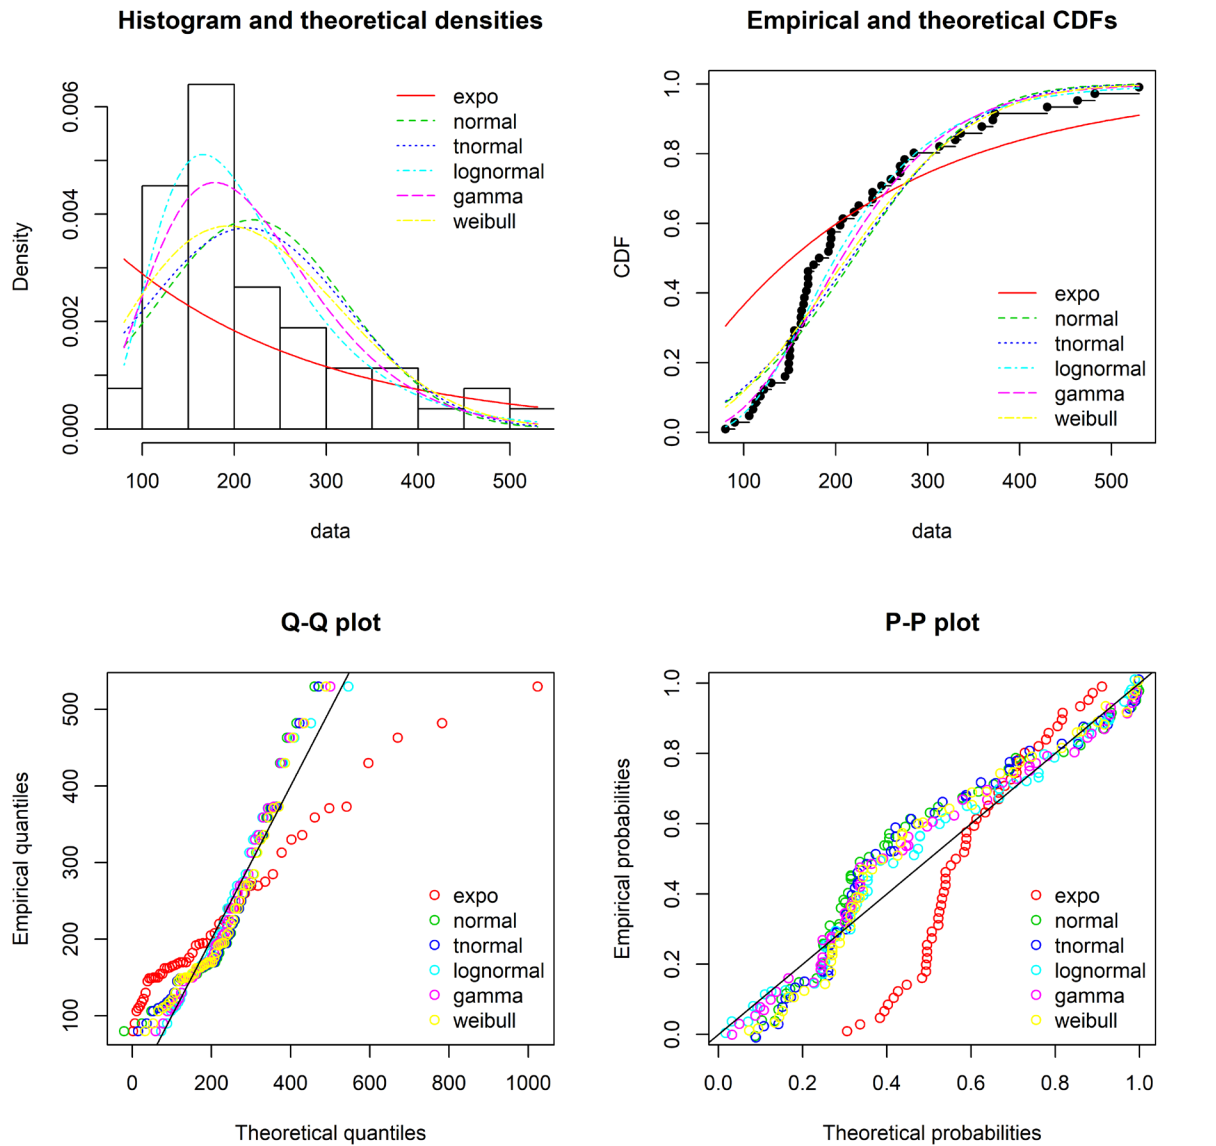
	expo	normal	tnormal	lognormal	gamma	weibull
Kolmogorov-Smirnov statistic	0.345	0.180	0.165	0.118	0.136	0.149
Cramer-von Mises statistic	1.675	0.401	0.348	0.112	0.186	0.278
Anderson-Darling statistic	8.419	2.254	2.001	0.584	0.988	1.578
Akaike's Information Criterion	680	645	643	626	630	637
Bayesian Information Criterion	681	649	647	630	634	641

eCDF comparison, with bootstrapped CIs

Cumulative probability

Duration in service point (seconds)
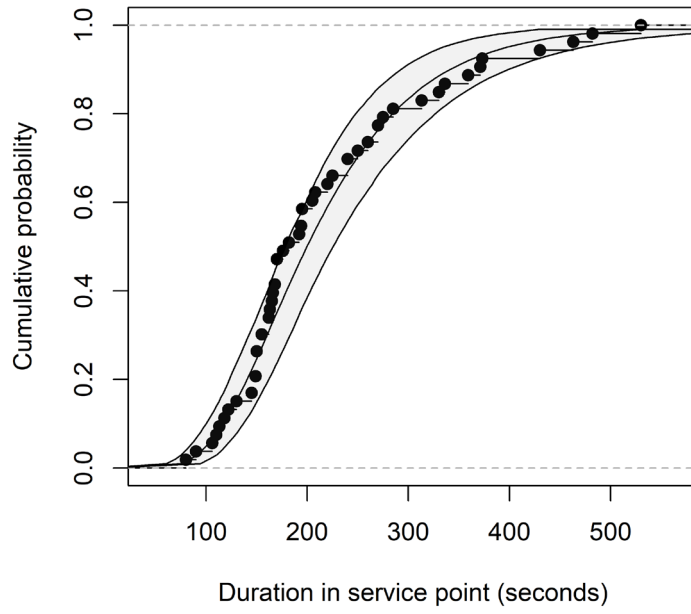
Appendix A.3 Immunisation (S3)

Histogram and theoretical densities

Empirical and theoretical CDFs

Q-Q plot

P-P plot
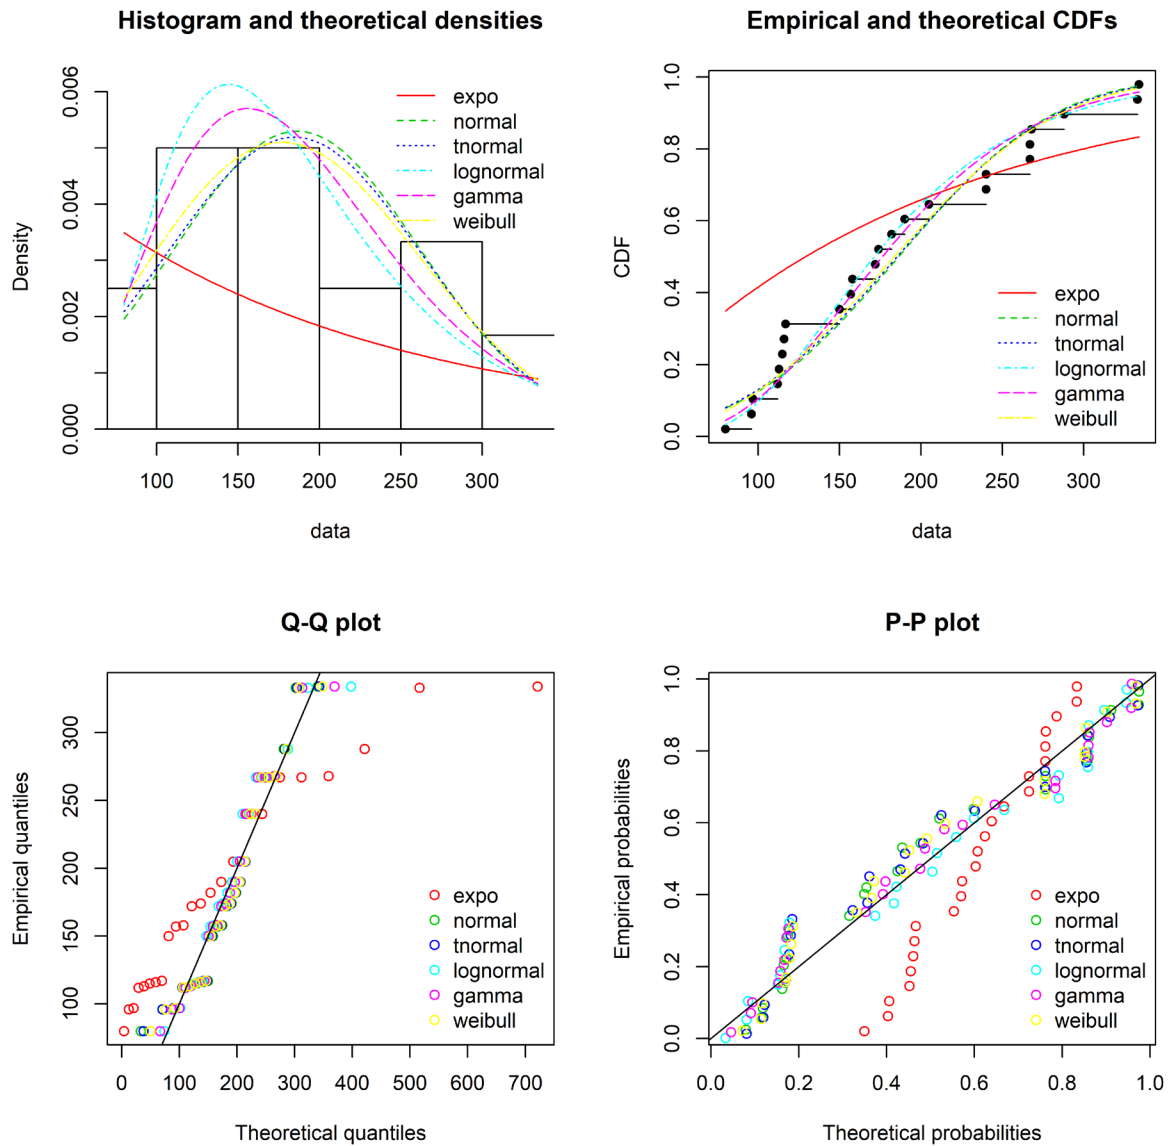
	expo	normal	tnormal	lognormal	gamma	weibull
Kolmogorov-Smirnov statistic	0.361	0.155	0.148	0.155	0.157	0.147
Cramer-von Mises statistic	0.770	0.097	0.087	0.065	0.066	0.075
Anderson-Darling statistic	3.969	0.628	0.574	0.441	0.453	0.496
Akaike's Information Criterion	301	280	279	277	277	278
Bayesian Information Criterion	302	282	282	279	279	280

eCDF comparison, with bootstrapped CIs

Duration in service point (seconds)