A complex, high-performance agent-based model used to explore tuberculosis and COVID-19 case-finding interventions

Marcus Low

Dissertation presented for the degree of Doctor of Philosophy in the Department of Computer Science, University of Cape Town

September 2023

Supervisor: Associate Professor Michelle Kuttel

Co-supervisor: Dr Nathan Geffen

I know the meaning of plagiarism and declare that all of the work in the dissertation, save for that which is properly acknowledged, is my own.

Contents

Acknowledgements

Associate Professor Michelle Kuttel assisted in preparing graphs. Dr Nathan Geffen helped with coding the model. Both Kuttel and Geffen gave feedback on various drafts of this dissertation. Janita Low edited the text of the dissertation.

Thanks

Thank you to Michele and Nathan for the unwavering support, encouragement, and guidance over the years. I am very grateful to both of you. Thank you to Janita for her love and support – and to her and Scott for putting up with the many hours I spent behind my computer.

Abstract

Tuberculosis (TB) claimed an estimated 1.5 million lives in 2021 and the COVID-19 pandemic resulted in 14.9 million excess deaths in 2020 and 2021 combined. With both these infectious diseases substantial pathogen transmission takes place before people report to health facilities. Diagnosing more people more quickly and placing them on treatment and/or isolating them is thus critical to achieving epidemiological control. Early diagnosis interventions include contact tracing and isolation, testing of asymptomatic people thought to be at high risk of infection, and in the case of TB, screening using chest X-rays. The impact of several such early diagnosis interventions on case detection have been studied in clinical trials, but the longer-term impact of these interventions on infections (incidence) and deaths (mortality) is not known. There are also unanswered questions as to the impact hypothetical future TB tests, for example allowing for more frequent testing, may have on TB incidence and mortality.

We developed an agent-based model (ABM) called ABM Spaces and used it to ask: (i) What is the impact of four different TB case-finding interventions on TB detection rates, incidence and mortality? (ii) What is the impact of test frequency and test sensitivity on tuberculosis incidence? And (iii) What is the impact of contact tracing and isolation and variable test turnaround times on COVID-19 diagnosis and mortality? Such agent-based modelling, in which disease transmission and progression is modelled at the level of discrete individuals, has increasingly been used in recent decades to model infectious disease interventions. Relatively few ABMs in the literature contain substantial social structure (for example associating agents with specific households, workplaces, and school classes). We illustrate that such ABMs with substantial social structure can be developed in a way that is epidemiologically sound and show that this type of ABM is well-suited to the modelling of social interventions such as contact tracing.

In the ABM Spaces model we found that testing of people at considerable risk of TB has a greater impact on TB incidence and mortality than mass X-ray screening, that the impact of the two interventions is additive, and that the impact of annual testing of high TB risk individuals is highly sensitive to HIV prevalence. We found that the relationship between test frequency and TB mortality and incidence is non-linear, with an inflection point at around the four-month mark. The COVID-19 version of ABM Spaces confirmed the potential of contact tracing and isolation to reduce incidence and mortality, but the effect was highly sensitive to test turnaround times.

Chapter 1 – Introduction

The world contains many complex systems. Some systems and causal relations can be studied using scientific experiments. In some cases, however, experiments are not feasible for practical or ethical reasons. In such instances models may at times be useful, providing they accurately capture the dynamics of interest. As argued by Vynnycky and White in their seminal handbook *An Introduction to Infectious Disease Modelling*, models are "just a simplified representation of a complex phenomenon". (1)

While this dissertation will focus on one highly specific type of modelling, agent-based models (ABMs) of infectious disease interventions, it is worth keeping in mind that ABMs are just one of many types of models. When the safety of medicines is tested in mice prior to humans, we are making use of a mouse model. A map of the London underground is a simplified model of the actual London underground. A series of ordinary differential equations can constitute a model of how a pathogen spreads in a population.

Epidemiological models are models designed specifically to explore epidemics, or the spread of diseases in populations. The purpose of such models can vary widely. Amongst others, models can be used to project the potential future trajectories of an epidemic or outbreak, to estimate the size of an epidemic, to better understand the factors impacting an epidemic, or to estimate the impact of different preventative measures. Models have, for example, been used to estimate the level of vaccination coverage required to prevent outbreaks of polio or measles (2) and to estimate the scale of South Africa's HIV epidemic (3). The nature and design of epidemiological models also differs widely – some models can be as simple as three or four ordinary differential equations that capture key transmission dynamics, while others can run to thousands of lines of computer code that simulate each instance of a pathogen being transmitted in a population.

In deciding on which model to use for a particular problem, a key question is whether the choice and configuration of the model simplifies the complex reality being studied in a way that is sound, meaningful and useful. While this is a question that has to be asked about every model individually, advances over the last century have meant that we now have excellent mathematical and statistical tools that we can use in these assessments. In the field of infectious disease modelling in particular, we have a set of solid foundations developed over decades of systematic study. Even so, it is not always obvious which choices about model design are the correct ones, and the strength of a particular set of choices still depends largely on the case the modeller can make in support of those choices. Often there are multiple valid modelling approaches with which to explore an epidemiological question.

Problem statement

One area of epidemiological modelling that is of particular interest to public health policy decisions is the modelling of diagnostic interventions. Such models can help us understand which specific diagnostic intervention, or combination of interventions, might be most effective. They might also elucidate relevant epidemiological dynamics, for example shedding light on why an intervention may be highly effective in one population and only moderately effective in another.

Diagnosing more people more quickly is of critical importance for the control of several infectious disease epidemics. For example, a person with untreated TB disease is estimated to transmit the bacteria to three to 10 people per year. (4) Each of those three to 10 people can in turn transmit to a further three to 10 people, and so on. Even though there is significant uncertainty about the accuracy of the estimate of three to 10 (apart from it being a wide range in itself), the underlying dynamic of substantial TB transmission prior to diagnosis and treatment is well established. With TB, as with several other diseases, earlier detection followed by isolation or treatment that makes someone non-infectious can thus have a substantial impact on onward transmission. What the optimal strategy is for early detection is however often not known.

Below we describe three specific problems, or epidemiological questions, relating to the diagnosis of TB and COVID-19. After these summaries of the epidemiological problems, we turn to the question of how to optimally model early detection interventions, given the inherently social nature of interventions such as contact tracing.

TB and COVID-19 make for particularly interesting infectious diseases in which to conduct this type of modelling given that they in some respects represent two extremes. TB is a very slow-moving infection with symptoms often only developing many months or years after someone inhales TB bacteria and treatment taking many months to complete. By contrast, people who contract the SARS-CoV-2 virus often develop COVID-19 symptoms within a few days and then typically recover or die within a few days or weeks after that. These differences in the “speed” of these infections have several implications for the strategies used to diagnose more people more quickly.

The problem of establishing how to diagnose more people with TB more quickly

Mycobacterium tuberculosis (TB) is a bacterial infection transmitted through the air. (5) The World Health Organization estimates that in 2020 TB claimed over 1.5 million lives and there were around 9.9 million cases of active TB worldwide. Of these only 5.8 million were diagnosed – leaving about 4.1 million people with active TB that were not diagnosed. (6) Figures for 2020 and the two subsequent years were impacted by a reduction in TB diagnosis due to the COVID-19 pandemic and the recovery from the pandemic – in 2022 there was an estimated 3.1 million undiagnosed cases of active TB.(7) Finding these three to four million “missing” patients is a key pillar of TB strategies worldwide.

Finding the missing cases involves both diagnosing more people with TB and diagnosing people earlier in their disease course. Though estimates vary widely, some have calculated that one person

with TB can transmit the bacteria to three to 10 others per year if not diagnosed and placed on treatment. (4) Earlier diagnoses thus offers individual level benefits – people do not become as sick before being treated – and population level benefits – people start treatment earlier on and are thus infectious for a shorter period of time and accordingly transmit the bacterium to fewer people. While it is relatively widely held that people with active TB become non-infectious within two weeks of starting treatment, a recent systematic review suggests that this is not the case and that 12% or 41% (depending on how it is tested for) of people still produced viable bacteria after two months. (8) Either way, early diagnosis and treatment initiation appears to have great potential for reducing TB transmission since people can remain sick and infectious for much longer in the absence of treatment, with some estimating average infectious periods of over a year. (5)

Traditionally, the approach to diagnosing TB has largely been to wait for symptomatic patients to present at healthcare facilities (called passive case-finding). As pointed out in a systematic review published in The Lancet in 2021: “Because tuberculosis care and prevention interventions that rely primarily on passive case detection and health facility-based screening strategies have insufficiently reduced tuberculosis incidence, many national tuberculosis programmes have promoted community-based active case-finding interventions.”(9) Active-case finding encompasses a wide range of interventions, with the common thread of not waiting for symptomatic patients to present at healthcare facilities. The systematic review found that "there is mixed evidence that active case-finding is effective at initially increasing tuberculosis detection when measured by case notification rates, and that active case-finding could reduce community prevalence of tuberculosis if delivered with sufficient intensity and coverage”.(9)

Though various aspects of the natural history of TB remains uncertain, significant progress has been made in recent years regarding our understanding of infection and progression to active disease. The traditional conception of TB as either latent or active, has given way to a more nuanced understanding whereby disease occurs on a spectrum ranging from latent, to insipient, to sub-clinical, to active disease. (9,10) This more nuanced understanding has implications for TB detection. For example, knowing that someone has latent TB has little predictive value as to whether that person will progress to active TB, and if so, when. Only around 5% to 15% of people with latent infection progress to active TB disease. (5) By contrast, someone whose disease has progressed to a sub-clinical stage is at a much higher risk of developing active disease, even though still asymptomatic.

Accordingly, the development of tests that can detect TB while it is still in a sub-clinical/asymptomatic phase has become an area of research. For example, there has been encouraging work on tests that predict progression to active disease by analysing a set of blood markers. (12) Such tests are however still under development and at best a few years from commercialisation.

There have also been promising signs that older technology in the form of chest X-rays can help with early detection. In South Africa’s first National TB Prevalence Survey, conducted in 2017-2019 with findings made public in 2021, of 234 confirmed cases of TB in the survey of over 35,000 people, 135 had chest X-rays suggestive of TB without having reported any TB symptoms. (13) Accordingly, a

number of studies are testing different strategies for the use of chest X-rays to aid early detection of TB.

Another early detection approach does not make use of a new test, but simply makes greater use of existing tests. One recent study showed encouraging results with a Targeted Universal Testing (TUT) strategy whereby people deemed at high risk of active TB were tested using the gold standard Xpert MTB-RIF molecular test – in this study high risk meant (i) living with HIV, or (ii) close contacts of someone with TB in the past year, or (iii) prior TB in the past two years. (14) In addition, more aggressive testing and X-ray screening strategies can be combined with different levels of contact tracing, for example, by offering X-ray screening or Xpert MTB-RIF testing to all household contacts of a person diagnosed with TB.

Some of the above strategies have been evaluated, typically in the form of cluster randomised trials in which entire communities or clinics are randomised rather than individuals. Such trials can generate information on the impact of interventions on case detection – but are generally too short in duration to meaningfully measure impact on incidence. In the above-mentioned trial of TUT, for example, 62 different clinics were randomised to either standard of care or TUT. The researchers found that TUT increased case detection by 17% compared to standard of care. (14) Findings from ACT3 and ZAMSTAR, two landmark trials of early case-finding interventions for TB, were mixed.(15) (16)

Cluster randomised trials are large, expensive, and typically only test a few interventions over a limited time period. Given differences between real world communities, there is a significant risk of confounding in cluster randomised trials, although this can to some extent be controlled for. The generalisability of findings from cluster randomised trials to communities different from those studied is also limited. Such trials may also not be set up, or continue for long enough, to detect the community-level impact over time of diagnosing people who would have been diagnosed anyway at an earlier stage in their disease course. Thus, though an essential research tool, cluster randomised trials of diagnostic interventions will inevitably leave at least some important epidemiological questions unanswered.

There remain broadly two gaps in our knowledge of the impact of new case detection strategies in the context of recent developments of our understanding of TB. Firstly, there are interventions, or combinations of interventions, which remain untested in randomised controlled trials. Since there are a large number of such potential interventions, and combinations of interventions, careful modelling could help decide which ones to advance to clinical trials. In addition, even though evidence from trials is essential to programmatic planning, it is unlikely that trials will provide all the answers given the time and expense involved. Therefore, some form of modelling based on the best available evidence will likely always have a role. Secondly, while clinical trials will yield information on case detection, few will continue for long enough, or be sufficiently powered, to provide reliable information on the secondary impact of such strategies on TB incidence. This is important since a key benefit of early detection is the reduction of onward infection. Again, modelling may elucidate the epidemiological dynamics at work and the impact of various interventions, not only on case detection,

but also on the longer-term impact on TB infections and ultimately mortality. Such work could, for example, be used to inform long-term cost-effectiveness analysis that takes into account the indirect consequences of increased case detection.

We address both these knowledge gaps by modelling a series of interventions and combinations of interventions and exploring the impact of these interventions on case detection, incidence, and mortality. Our work thus attempts to shed light on the real and urgent question as to what early detection interventions should be prioritised in TB programmes in contexts such as those in South Africa.

The problem of defining optimal target product profiles for TB diagnostics

While the optimal use of existing TB case-finding interventions is one unsolved problem, another is to pinpoint which characteristics developers of new TB diagnostics should focus on. This is in part a matter of expert consensus, as for example shown in a World Health Organization (WHO) report on target product profiles for TB diagnostics,(17) but it is also a question that should ideally be explored using the best available evidence. Since the available evidence does not always provide clear answers to important questions, such as the optimal frequency of TB testing, there is a potential role here for modelling. For example, based on our existing knowledge of TB and TB diagnosis, models could be developed that simulate the impact of different levels of test frequency. The outputs from these models could then be taken into account in the development of target product profiles for TB diagnostics. (Note: While the assessment of “optimal” strategies should ultimately take into account factors such as cost-effectiveness and acceptability, the focus of the modelling presented in this dissertation is on reducing infections and deaths.)

That new and better TB diagnostics are needed is widely acknowledged by the WHO and other leading TB institutions such as the Stop TB Project. Current gold standard molecular tests, such as the Xpert MTB/RIF, require people to produce sputum – something many people find hard to do. A urine or saliva-based test would be a significant step forward, providing it has sufficiently high sensitivity and specificity (existing urine tests only have high sensitivity in people living with HIV who have severely compromised immune systems). Though point-of-care devices are under development, the currently available molecular tests are generally still processed in laboratories. Point-of-care, or ideally home tests, would both make testing more convenient and facilitate more frequent testing. Either way, the requirement for sputum and the need for laboratories with the current tests, make high frequency testing unfeasible. The most ambitious TB testing programmes we are aware of, such as TUT, offer annual tests only to certain sub-groups considered to be at high risk of TB. We are not aware of any programmes offering TB tests more frequently than once a year.

More frequent testing can benefit people with TB through earlier diagnosis and treatment, thus preventing symptomatic disease or shortening its duration. Apart from the individual benefits, more frequent testing can also have public health or population level benefits by reducing TB transmission.

Just how more frequent testing might reduce TB transmission is, however, not well understood. Accordingly, by modelling the impact of several levels of test frequency and test sensitivity on TB incidence and mortality, we shed light on the question of how often we would have to test for testing to result in substantial reductions in TB incidence. The cost and acceptability of such higher frequency testing is however beyond the scope of this modelling exercise.

The problem of establishing how to diagnose more people infected with SARS-CoV-2 more quickly

Coronavirus Disease 2019 (Covid-19) is an infectious disease caused by the SARS-CoV-2 virus. After first being detected in Wuhan, China at the end of 2019, the virus rapidly spread around the world causing a global pandemic claiming millions of lives in the subsequent two years.

By mid-2020, many countries had implemented extended and far-reaching lockdowns in response to the pandemic, often at significant economic cost. A key question in the early months of the epidemic was what non-pharmaceutical interventions (NPIs) governments could implement in order to exit such lockdowns as safely as possible. (18) Some potential NPIs include promotion of hand-washing and other hygiene measures, promotion of or legal requirement of mask-wearing in public, various physical distancing regulations, border closures, school closures, and contact tracing and isolation (CTI). (19)

CTI involves tracing the recent contacts of a person who tested positive for SARS-CoV-2 and isolating these contacts, and in some cases also testing contacts. Teams of tracers traditionally conduct CTI, but it can also be done through, or assisted by, various mobile phone applications. WHO has endorsed CTI as a key strategy to slow transmission of SARS-CoV-2. (20) CTI has been implemented in a wide range of countries – maybe most notably in China, Singapore and South Korea. (18-20) One Lancet study concluded that “each country should have an effective find, test, trace, isolate, and support system in place”. (18) Isolation refers to the movement of people with the infection being restricted, voluntarily or legally, while quarantine refers to the movement of people at risk of infection being restricted, voluntarily or legally. The work described here does not differentiate between these distinctions. We therefore use isolation to refer interchangeably to isolation or quarantine.

Since there are economic, human resource and opportunity costs related to CTI, it is important to understand the effectiveness of such programmes and the factors impacting their effectiveness. The efficacy of CTI programmes in preventing infections and mortality will be different in different settings and will depend on factors such as test turnaround times (TaT) and on whether contacts are tested or not. TaT varied substantially in South Africa in 2020, at times exceeding seven days. It seems plausible that, while potentially effective with a TaT of two days, CTI may become ineffective with a TaT of seven or more days. In our work we thus explore the problem of quantifying the impact of longer test turnaround times on the efficacy of CTI.

The problem of modelling early detection interventions in complex social networks

In intervention models, a first objective is typically to develop a model that recreates key dynamics of how the disease in question is transmitted in real-world populations. If a sufficiently accurate transmission model can be developed, the model can then be adapted to see how different screening or testing strategies might improve diagnosis compared to a baseline and what impact such improved diagnosis might have on the trajectory of the epidemic in question.

An intervention such as contact tracing can be modelled in a wide variety of ways. In its simplest form it may be assumed that contact tracing increases diagnosis in an entire population by a certain factor and in a model based on ordinary differential equations the equations can be adjusted to include this additional factor. At the other end of the spectrum, in an ABM, where modelling happens at the level of the individual and every instance of disease transmission is discretely modelled, contact tracing can be modelled as the process of tracing the actual contacts of every diagnosed individual in the model.

A good rule of thumb in the modelling world is that one should use the simplest possible model that adequately captures the key dynamics of interest. There are instances however in which the complexities of the phenomenon being studied are such that certain epidemiological dynamics can only be explored using relatively complex models.

When modelling interventions on TB detection and transmission dynamics, we propose that the choice of model should be guided by three factors: (i) Whether the type of model captures the known epidemiological dynamics of interest, (ii) whether relevant data exists with which to inform and calibrate the model, and (iii) whether the model design allows for the meaningful modelling of interventions.

Starting with the epidemiological dynamics of interest, the spread of TB in a community (similar dynamics apply to COVID-19) can be thought of in terms of two processes. Firstly, there is the natural history of TB, or the way TB manifests in particular individuals. There is significant heterogeneity in this regard, with, for example, not everyone exposed to the bacterium becoming sick with TB. Secondly, there is the process by which people encounter each other and the bacterium is transmitted. In the real world, this is also a highly heterogeneous process with people meeting at schools, workplaces, in households, and on public transport. In the first three of these settings, agents mix with the same group of agents on every timestep and in the last, agents can mix with different agents on different timesteps.

Modelling of TB and COVID-19 testing or screening interventions should ideally capture both these natural disease processes and the transmission processes in sufficient detail. In both cases models should preferably contain the type of heterogeneity seen in the real world. For example, if we consider that an X-ray may detect TB in a specific person while that person is still asymptomatic, then we'd

also be interested in the extent to which this early detection results in the person being infectious for a shorter period of time, and accordingly transmitting TB to fewer people than otherwise would have been the case. In this way, the early detection of TB in a worker may prevent an outbreak in that worker's workplace and home, which in turn can prevent further transmission in the school class attended by members of the worker's household. Similarly, TB infection can also spread through these networks in a comparable manner. Intervention modelling should ideally capture these transmission and prevention patterns in sufficient detail to allow for the development of emergent phenomena, rather than simply assuming certain rates and discounting the possibility of any unexpected dynamics emerging from the complexity of the network.

One way in which such complex transmission dynamics could be modelled is in the form of ABMs that contain significant social network structure. In such models each person exists as a discrete agent. The disease process for each agent exposed to a pathogen can be modelled to reflect published estimates of the natural history of the disease in question, with heterogeneity built in by sampling from plausible probability distributions rather than assigning the same parameters to all agents. These agents can then be made to interact within a structured model space in which agents are associated with specific households and school classes and/or workplaces. In this way the two key processes of disease progression and transmission in social settings can be modelled simultaneously and in an interconnected way.

Regarding the available data, the natural history of TB has been widely studied, although some key parameters remain uncertain. This is a problem for TB modelling no matter the type of model – although models that take into account the uncertainty arguably provide more realistic estimates (or ranges of estimates). Though the situation with COVID-19 is somewhat fluid due to the development of new variants and the impact of vaccination and immunity acquired from infection, the natural history of COVID-19 has nevertheless been studied extensively and sufficient data has been published with which to parameterise and calibrate this type of model.

Regarding transmission dynamics there is an emerging evidence base that can be drawn on exploring the types of settings in which TB transmission is most likely to take place. Such estimates can inform ABMs in various ways – ABMs can, for example, be calibrated to reproduce transmission patterns similar to those observed in the published literature. Since SARS-CoV-2 is also largely transmitted through the air, data on "shared air" (a measure of the extent to which people are breathing the same air based on in-door CO2 levels) from TB studies can be used to inform COVID-19 models.

Finally, since ABMs operate at the agent level, they allow for the direct implementation of interventions. For example, when an agent is X-rayed as part of a mass X-ray campaign, the sensitivity of the X-ray screening is applied to that specific agent and there is a specific and discrete test result. This is different from models where screening tests are dealt with in the aggregate. When such ABMs also contain social structure in the form of households, school classes, and so on, they allow for interventions to occur within realistic networks of agents – thus, as with the modelling of disease itself, interventions are also modelled with more context and more heterogeneity than would

be the case in simpler models. ABMs that contain substantial social structure are thus well-suited to capturing the complex consequences that interventions may have on populations.

A similar argument can be made for modelling the impact of CTI and different test turnaround times on COVID-19 incidence and mortality using ABMs with social structure. In general, many published SARS-CoV-2 models are compartmental models rather than ABMs (see for example (19, 20, 21). Compartmental models (discussed in more detail in chapter 2) describe the movement of people between different compartments (typically susceptible, infectious, recovered) using ordinary differential equations. Though significant complexity can be built into such models using various mathematical techniques and by adding compartments, they do not model at the agent level, but at the level of groups or compartments – which results in homogenous mixing. This arguably makes them less appropriate for modelling interventions such as CTI, that by its nature takes place within complex social networks that are not typically captured in compartmental models.

In contrast to compartmental models, ABMs with realistic social structure can provide a way to explore the complex network effects and consequences of interventions such as CTI on the transmission of SARS-CoV-2. By modelling at the agent level, where each agent represents an individual, such models better account for heterogeneity within populations than is the case with compartmental or population-level models that by definition model groups rather than individuals. By including structure, such ABMs can factor in the complex dynamics whereby some agents within a household may be employed, others might go to school, and others may be unemployed – dynamics that at scale could lead to emergent effects such as early diagnosis and isolation in schools and workplaces potentially shielding older unemployed persons at home from infection. In a recent review, four of seven models investigating CTI for the control of earlier outbreaks of Severe Acute Respiratory Syndrome and/or Middle-Eastern Respiratory Syndrome were ABMs(26). A review of the literature (see chapter 2), however, suggests that ABMs with substantial social structure (households, schools, workplaces, and so on) is an under-explored area of epidemiological modelling.

ABMs with substantial social structure pose at least two significant challenges. Firstly, they are harder to parameterise and to calibrate since they tend to have more input variables than simpler models. There is thus a greater risk of the “garbage in, garbage out” problem – though this can also be a risk with very simple models. More parameters also make it harder to ensure that variables impact each other in ways that are sufficiently realistic and relevant to the question being explored. Thus, even if it is not a case of “garbage in”, the output might still be “garbage” if the model fails to capture the relationship between variables with sufficient accuracy – something that becomes more likely as models get more complex. This issue with parameterisation and calibration is best thought of as an epidemiological problem rather than a computing problem.

Secondly, the more complex models become the more computing power is generally needed to run them. Simulating larger populations, allowing more agent actions per time step, and modelling longer time horizons also tend to increase computational burden. Improvements in computing power and affordability over the last two decades have resulted in an increase of complex models such as ABMs

in the academic literature. But despite these advances, computational burden remains a limiting factor in the design of ABMs described in the literature (as we will show in chapter 2). That is to say some decisions in model design are not driven purely by epidemiological or scientific considerations, but by questions such as how long it will take to run the model. Part of the problem here is that personal computers are not yet fast enough, and access to more powerful computing resources not yet universal enough, to make computational burden irrelevant to model design. Another part of the problem is that model designers typically use relatively high-level modelling packages that, while making it easier to design models, do not offer the high execution speed and configurability of compiled programming languages such as C++. This issue is best thought of as a computing problem, rather than an epidemiological problem.

Such questions of complex versus simple models can up to a point be dealt with in the abstract, but there is also value in simply building and parameterising complex models with as much scientific rigour as possible, seeing what we can learn from the process and subjecting the resulting models to the scrutiny of the scientific community. This is the approach we've taken here.

In short, although ABMs have increasingly been used to model infectious disease interventions, they have generally, and especially in relation to TB, not included the type of substantial social structure that may be key to accurately capturing important nuances of how pathogens spread in populations. ABMs with substantial social structure thus represent an underexplored solution to the problem of how to optimally model infectious disease interventions with an inherently social component such as contact tracing.

Aims of research and research question

The research presented here straddles the intersection of epidemiological modelling and applied computer science (or scientific computing). Accordingly, our research aims and research questions contain elements of both fields and does not neatly fit into either.

The principle aim of this work was to develop a complex ABM with substantial social structure and to use that ABM to explore the impact of several diagnostic interventions aimed at earlier detection of TB and COVID-19.

In support of this over-arching aim, we set three objectives:

i. To conceptualise and design an ABM with substantial social structure in such a way that it would allow for the modelling of TB and COVID-19 case-finding interventions.
ii. To efficiently code this model so as to allow for thousands of model runs to be executed on a personal computer in minutes, and
iii. To use the model to explore three real-world epidemiological questions.

The three real-world epidemiological questions we explored are:

i. What is the impact of four different TB case-finding interventions on TB detection rates, incidence and mortality?
ii. What is the impact of varying levels of test frequency and test sensitivity on TB incidence and mortality? And
iii. What is the impact of contact tracing and isolation and variable test turnaround times on COVID-19 diagnosis and mortality?

Approach and contribution

To answer the above questions, we developed a model called ABM Spaces. In brief, ABM Spaces is an ABM with social structure in the form of households, workplaces, school classes, taxis, and city blocks. It allows for the modelling of several different types of diagnostic interventions within a heterogeneous population with realistic social structure. It is coded in C++ and hundreds of simulations can be run in minutes on a standard laptop computer. It is calibrated using, among others, previously published estimations of tuberculosis transmission patterns in a high-density township in South Africa. (See chapter 3 for a detailed discussion of ABM Spaces.)

ABM Spaces is freely available on GitHub and can be relatively easily modified by other researchers to explore similar epidemiological questions to those presented in this dissertation. ABM Spaces is thus both of value as scientific software that other researchers can use and build upon and as a proof of concept as to how this specific type of modelling can be done in a compiled language such as C++.

The aim, however, is not only to develop scientific software, but to show how the software can be used in a scientifically rigorous way to conduct complex agent-based modelling. Our contribution here is to present an epidemiological methodology for developing, parameterising, and calibrating ABMs with substantial social structure. There are relatively few ABMs of this nature in the published literature and some aspects of our approach to this type of model are, as far as we can establish, novel. The aim was also to illustrate that this type of modelling work can be relevant to specific real-world policy questions. Thus, the chapters on TB and COVID-19 diagnostic interventions are aimed at answering actual policy questions. Much attention was accordingly paid to reviewing the relative TB and COVID-19 literature and parameterising and calibrating both the TB and COVID-19 versions of ABM Spaces and ensuring the model design was appropriate to the questions being explored. As a result, the work presented here also contributes to the literature on epidemiological questions relating to the optimal diagnostic strategies for TB and COVID-19.

Our chapter on case-finding interventions for TB is novel, firstly, in that it includes sub-clinical TB – a review article published in 2018 stated that “Most models of TB transmission have not included a subclinical TB stage, which may constitute half of the average total duration of active TB prior to treatment.”(10) Secondly, our TB model is one of the first to make use of findings from South Africa’s first National TB Prevalence Survey (published in 2021). (13) Thirdly, this is one of very few models we have found in the literature that makes use of an agent-based modelling framework with realistic social structure to explore the impact of different interventions on TB transmission. It is also the first study, to our knowledge, which evaluates the impact of contact tracing, mass X-ray screening, and TUT in a single model and as such provides policy-makers with useful background to consider when deciding on which early detection strategies to implement.

Our chapter on TB test frequency and test sensitivity is novel in that it models test frequency in the context of an ABM with social structure and with a sub-clinical TB phase. It identifies that the relationship between test frequency and TB incidence is not linear and that there is a significant inflection point at around the four-month mark. Our work suggests that the development of a TB test that can be administered every three months, or ideally even more frequently, should be a high research priority. Such frequent administration would require the test to be cheaper and much easier to administer than currently available TB diagnostics.

Our chapter on COVID-19 diagnosis adds to previously published work exploring the impact of CTI and TaT using different methodologies. For example, one branch process modelling study found that, in most scenarios, highly effective contact tracing and case isolation is enough to control a new SARS-CoV-2 outbreak within three months, but that “the probability of control decreases with long delays from symptom onset to isolation”. (27) Such delays could, for example, be due to long TaTs. This branch process model did not contain any compartments or social structure – but merely looked at how many people would be infected by each infected person (the branching), and how that process can be controlled with CTI. A similar branch-process modelling study found that “minimising testing delay had the largest impact on reducing onward transmissions” and suggested that test turnaround times (TaTs) of less than three days were required to bring the effective reproductive number under 1.

(28) A model described as a “model of individual-level transmission stratified by setting” found that CTI in combination with other interventions may contribute to substantial reductions in SARS-CoV-2 transmission. (29) While this model included transmission from infected individuals to individuals in households, schools and workplaces (based on probabilities derived from real-world data), these settings were not persistent in the model, and thus did not allow for the full network effect of transmission chains to be explored. Our findings, first published as a preprint in late 2020, added to the growing literature on the efficacy of CTI and the importance of test turnaround times for the control of COVID-19.

Structure of the dissertation

This dissertation is structured as follows:

In chapter 2 we contrast ABMs with compartmental models (the type of model most widely used for infectious diseases). We then provide a brief overview of the different types of ABMs and considerations in the design of ABMs. That is followed by a literature review of use of ABMs to explore questions relating to first TB and then COVID-19 interventions. We focus on the types of ABMs used, the environments in which agents interact, the complexity of the models, and other key design choices and trends.

In chapter 3 (methodology) we provide a detailed description of how the ABM Spaces model works, and we explain the reasoning behind the various design decisions and why we consider ABM Spaces to be an appropriate model with which to explore questions relating to early TB detection interventions. We also provide a detailed explanation of how both the TB and COVID-19 versions of ABM Spaces were parameterised and calibrated. In addition, we discuss the coding of ABM Spaces.

In chapter 4 we present the findings of an implementation of ABM Spaces used to explore the impact of four TB case-finding interventions on TB detection rates, incidence and mortality. Here we attempt to make a contribution to a real and urgent health policy question as to which interventions will most effectively diagnose more people with TB more quickly. The analysis in this chapter illustrates our core contention that relatively complex agent-based models can meaningfully be used to explore real-world epidemiological questions and that by coding in a compiled language such as C++, such models can be sped up significantly to take only minutes.

In chapter 5 we use a lightly modified version of ABM Spaces to explore the impact of variable test frequency and test sensitivity on TB incidence and mortality. Our analysis here can help inform work on defining target product profiles for TB diagnostics.

In chapter 6 we present the findings of an implementation of ABM Spaces used to explore the impact of contact tracing and test turnaround times on COVID-19 infections and mortality. While the core of the work presented in this dissertation relates to TB, ABM Spaces was first applied to COVID-19 as a

proof of concept (since modelling COVID-19 transmission in this type of model is simpler than TB) and since it became apparent that the model might usefully be used to explore certain real-world dynamics that were of public interest at the time the model was first being coded. Accordingly, an earlier version of this chapter was previously published as a preprint on medRxiv in October 2020 (having made the work available as a preprint, a decision was taken to focus on the TB chapters of this dissertation rather than pursuing journal publication). (30)

In chapter 7 we briefly recap some of our key findings and discuss what types of analysis can and cannot be done with agent-based models such as ABM Spaces.

Chapter 2 – Agent-based models and their use in modelling TB and COVID-19 interventions

In this chapter we start by contrasting agent-based models with compartmental models. We then provide a brief overview of the different types of ABMs and considerations in the design of ABMs. That is followed by a literature review of the use of ABMs to explore questions relating to diagnostic interventions for TB and COVID-19 respectively. We focus on the types of ABMs used, the environments in which agents interact, the complexity of the models, and other key design choices and trends.

How ABMs differ from compartmental models

Much modelling of infectious diseases is done using compartmental models along the lines of a model first published by Kermack and McKenrick in 1927. (31) The classic example of a compartmental model is an SIR model, in which people are allocated to compartments labelled “susceptible”, “infective/infectious” and “recovered/removed”. These models are also sometimes referred to as SEIR models when an “exposed” compartment is added. Designing compartmental models to have fewer or more compartments is relatively trivial and compartmental models often have many more than just three compartments. In compartmental models movement between compartments are defined by equations. As Vynnycky and White explain in their book *An Introduction to Infectious Disease Modelling*, these are typically difference or ordinary differential equations that take into account the size of compartments and factors like the contact rate between people, the risk of transmission per contact and the time period for which someone remains infectious. (1) Solving these equations is trivial using software like the R package deSolve and the computational burden is usually relatively low.

In their classic form compartmental models are deterministic – which is to say that given the same initial parameters they will always produce the same outcomes. Initial parameters may be varied to account for uncertainty regarding the values or to test sensitivity to different parameter values. A deterministic compartmental model might thus be run many times with varying parameters to produce a range of outputs.

One key characteristic of compartmental models is that everyone in a compartment is assumed to be subject to the same rates – for example, if people recover in five days, people might move from the infectious compartment to the recovered compartment at a rate of 1/5 per day. Accordingly, compartments can contain fractions – there may for example at a certain time point be 21.4 people who are in the infectious compartment. While this tendency toward treating people as homogenous groups makes compartmental models comparatively simple and easy to work with, it does mean that

compartmental models may not capture certain epidemiological dynamics that result from high heterogeneity between people, such as some people having more social contacts than others.

Compartmental models can be customised to introduce elements of randomness, referred to as stochasticity, and to ensure that only whole numbers are reflected in the various compartments. In practice, compartmental models are modified in a wide variety of ways, can contain elements of other types of models, can contain elements of heterogeneity, and become highly complex. An example of a complex compartmental model is the Thembisa model. It is used by the South African government and UNAIDS for HIV estimates in South Africa. (3) An influential COVID-19 model first published in March 2020 by researchers at Imperial College London was also a compartmental model – a version of this model was later published in the journal *Science*. (32)

Compartmental models are particularly useful in helping people to understand the possible trajectories that a disease outbreak might take and are often used for this purpose. One review of COVID-19 forecast models found that, "SEIR models are among the most widely adopted mathematical frameworks to describe disease dynamics and forecast potential contagion scenarios". (33) Such models are also used to elucidate the epidemiological dynamics at play in an outbreak or epidemic and to estimate the potential impact of interventions such as vaccination or quarantine.

As with any other model however, the usefulness and accuracy of compartmental models depend on the accuracy of the model parameters and availability of data for calibration – which in the case of COVID-19 was particularly uncertain in the early stages of the pandemic, while data used for the HIV model mentioned above are much more reliable. Typically, compartmental models have the benefit of relying on a relatively small number of parameters. They are thus comparatively simple to parameterise, as long as reliable parameter estimates are available.

While compartmental models remain the cornerstone of infectious disease modelling, there are many types of modelling being used – a review of COVID-19 prediction models identified machine learning models, regression models, and long short-term memory models, among others. (33) Our focus here is on agent-based models (ABMs). We contrast ABMs with compartmental models as a means to explain how they work, but also to lay the foundation for our later explanation of why we chose ABMs to address the research questions in chapters 4, 5, and 6.

ABMs – also sometimes called individual-based models or microsimulations – differ from compartmental models in that they model at the agent level rather than at the compartment level – which is to say processes in the model impact agents individually rather than as groups or compartments of agents. Vynnycky and White say of ABMs, which they call individual-based models: "These track what happens to each individual in the population and allow chance to determine whether or not they become infected and infectious at each time step (i.e. an individual-based approach)." (1) This "chance" element is why ABMs are classified as a form of stochastic modelling.

According to Railsback and Grimm, ABMs are:

> "models where individuals or agents are described as unique and autonomous entities that usually interact with each other and their environment locally. Agents may be organisms, humans, businesses, institutions, and any other entity that pursues a certain goal. Being unique implies that agents usually are different from each other in such characteristics as size, location, resource reserves, and history. Interacting locally means that agents usually do not interact with all other agents but only with their neighbours in geographic space or in some other kind of 'space' such as a network. Being autonomous implies that agents act independently of each other and pursue their own objectives."(34)

At a most basic level, we propose that an ABM is a model in which (i) every individual is associated with a unique ID and at least one state variable and (ii) in which at least some individuals' state variable(s) can be impacted by risk factors or by other individuals over time.

A key characteristic of ABMs is that they allow for modelling with much greater heterogeneity than compartmental models. For example, rather than subjecting everyone in a compartment to a certain rate that encapsulates the combined risk of infection, every single agent is individually exposed to a risk and it is calculated for each agent individually whether or not infection takes place. This property of exposing each agent to a risk, combined with the mediation of contact between agents through social structure can result in models that more closely capture the real-world dynamics of how disease spreads within social networks. This is why Vynnycky and White write, "One advantage that these models have over less structured deterministic models, is that they can be used to explore the impact of interventions targeted at the household, school or workplace."(1)

The ability to model in more detail and with more nuances comes at a cost. The number of calculations involved in running an ABM tends to be orders of magnitude more than in compartmental models. This is because on each timestep all or most of the agents in the model are potentially subject to unique calculations, however simple or repetitive those calculations may be. In compartmental models, by contrast, there is roughly an equal number of equations per timestep as the number of links between compartments in the model. It is possible that an SIR model may have only two equations to solve per timestep, even if it is modelling a population of 10,000 people – while an ABM modelling 10,000 people may in some manner have to consider each of those 10,000 individuals on each timestep.

ABMs may also require greater complexity in various other design aspects. While in a compartmental model contact between people can be abstracted to a rate defined by a differential equation, ABMs will typically function at a lower level of abstraction and let chance play a greater role in determining whether two agents encounter each other and whether transmission takes place. By modelling at a lower level of abstraction, more parameters may be needed and accurate model calibration may become harder. As stated by Vynnycky and White, "the obvious disadvantage of such highly structured models are that they are difficult to set up and they can be slow to run. A further challenge is that the most detailed models require many input parameters (some of which are not readily

available). As the number of input parameters increases, it becomes increasingly more difficult to elucidate the contribution of a given input parameter on the overall outcome, and thus many simulations may need to be run in order to obtain useful predictions."

Vynnycky and White point out that "with the development of fast computers, these disadvantages should become less of a problem in future." While faster computers are indeed transforming this field, we propose that there is also considerable progress to be made through better and more efficient scientific programming – an element that has arguably been neglected in the ABM literature. In fact, one of the key objectives of the work presented in this thesis is to illustrate that efficient coding of highly complex ABMs in a low-level language such as C++ allows for thousands of model runs to be executed in a matter of minutes on standard hardware.

The trade-offs between the complexity of ABMs and the relative simplicity of compartmental models has at times lead to academic disagreement. For example, in 2014, Williams wrote in response to an ABM-based paper by Hontelez and colleagues (35) that "their more complex models rely on unwarranted and unsubstantiated assumptions". (36)

Types of ABMs

ABMs can range widely in size, ranging from ABMs with only a handful of agents (e.g. a hospital ward) to ABMs with many thousands of agents (e.g. a suburb or city). In addition to the number of agents, ABMs can also differ widely in how much complexity is assigned to agents. In some models, agents may have only a few state variables each, while in more complex ABMs, each agent might be associated with tens or even hundreds of variables.

There are typically two types of processes that impact an agent's state variables in an ABM. Firstly, there are processes internal to agents – for example, the process by which an agent exposed to a virus may over time transition to becoming infectious and then recovered. These processes can be thought of as internal since, once initiated inside an agent, the process unfolds without reference to the world outside of the agent, unless the process is modified through a specific interaction with the outside world. There is significant variability in the complexity of such internal processes in ABMs. For example, some infectious disease ABMs may simply let all agents be infectious for a fixed number of days – while more complex models may build in multiple disease stages and determine the time spent in each stage by sampling from probability distributions based on the published literature on the natural history of the disease in question.

Secondly, there are processes by which agents encounter other agents or environmental conditions in the model world (We use the term "model" to refer to the entire model and we use the term "model world" to refer only to the environment or space in which agents in the model interact.). A susceptible agent might for example encounter an infectious agent and may or may not get infected – if infection does take place, it may then initiate an agent's internal disease process. Models differ both in how

they determine which agents encounter each other and in the mechanisms used to determine whether or not transmission takes place during these encounters.

Determining which agents encounter each other largely depends on the structure or shape of the model world. Some models, for example, make use of a linear space whereby all agents are situated along a line. In such models, an agent may for example only ever encounter agents that are within 10 positions of them on the line. Alternatively, agents could be distributed across a grid or lattice and only ever encounter their neighbouring agents. Agents can also be distributed over more complex two- or three-dimensional shapes and in more complex algorithmically determined networks. In some models, agents can interact randomly with other agents and with agents far away from them. Models also differ in whether all agents have the same number of social encounters or whether some agents have more than others. Rather than letting agents exist in a linear or multi-dimensional space, some ABMs create realistic social network structures by letting agents be associated with specific households, workplaces or school classes. ABMs can also be spatial – which means agents encounter each other in a two- or three-dimensional space analogous to a real-world setting.

There are several methods for determining whether or not transmission takes place during an encounter between agents. Most methods rely on randomly drawing a number and then checking whether that number falls into a certain range. For example, when a susceptible agent encounters an infectious agent and the risk of transmission during an encounter is set to 20%, a number might randomly be drawn from the range 0 to 1 – if the number drawn is less than 0.2, transmission takes place, and if it is greater than 0.2, transmission does not take place. While this example applies to direct encounters between agents (relevant e.g. to sexually transmitted infections), encounters can also be mediated by environments such as hospital wards, households, or workplaces. In such cases the risk of a susceptible agent in an environment contracting a disease is impacted not only by the infectiousness of the disease, but also by the number of infectious agents in the environment and the likelihood of encountering them. In its simplest form the risk of a susceptible agent in such an environment contracting a disease would be a certain infectiousness factor times the number of infectious agents in the environment (such a design may for example be relevant when modelling the transmission of airborne pathogens like tuberculosis in small, shared spaces where doubling the number of infectious agents can lead to double the amount of TB bacteria in the air). But since the risk of transmission of a pathogen in a ward may in many cases not in fact double if four as opposed to two people in a ward are infectious, Vynnycky and White propose instead using the Reed-Frost formula. The Reed Frost formula states that the risk of transmission is equal to $1 – (1 - p)^I$ – where I is the number of infectious agents in the environment and p indicates the risk of a contact between a susceptible and an infectious agent resulting in transmission. (1) Any number of other formulas could also be constructed. The choice of formula should ideally be guided by an assessment of how the pathogen being modelled is in fact transmitted in the real world.

ABMs also differ in the timesteps used in simulations – typically days, weeks, or months. While models are configurable, analyses reported in the published literature describe a certain number of model runs, each for a certain duration of time (time in the model world). In general, larger, more

complex models with shorter timesteps will run more slowly and may accordingly in some cases be run for shorter periods, with smaller numbers of agents, and with smaller numbers of model runs. As pointed out by Vynnycky and White, “the number of runs needs to be balanced against computational burden”. (1)

Finally, ABMs can be built or coded in many different ways. Some software packages like NetLogo (37) and AnyLogic (38) allows for ABMs to be designed using a combination of coding and use of a graphical user interface. Some programming languages like Python and Java have packages that provide tools with which to build ABMs. ABMs can also be developed in programming languages such as Python, Java, Julia and C++ without using ABM-specific packages. In deciding between these options, several trade-offs have to be considered. Using high-level packages may allow for ABMs to be developed quickly but will tend to be less customisable and to result in slower code. On the other hand, coding directly in a programming language will take longer, but will provide essentially unlimited customisability and probably faster code. In addition, developing an ABM in a compiled language like C++ is likely to take longer than doing so in a higher level, uncompiled language such as Python, but a model coded in C++ is likely to run much faster than one coded in Python. Either way, computational burden is not only a product of the model itself, but also of how efficiently and in what language the model is programmed.

How ABMs have been used to explore tuberculosis interventions

We conducted a literature review of how ABMs are being used to explore questions relating to TB interventions. We focus on the types of research questions that ABMs are used for, the model worlds in which agents interact, the complexity of the models, and other key design choices and trends. We searched PubMed and Google Scholar for the terms “tuberculosis” and “agent-based model” or “individual-based model” or “microsimulation”. We restricted search findings to peer-reviewed articles and sifted out articles that were not about TB interventions. We paid particular attention to articles published in leading medical journals. We included systematic reviews in our analysis. Our review is not systematic in the sense that we did not count and classify all models found in our search – rather than exact quantification, our purpose was to identify broad trends in the use of ABMs for modelling TB interventions.

Infectious disease ABMs have become more popular in recent years, although it is not clear whether TB ABMs in particular have become more popular. In a review of agent-based models applied to infectious disease interventions, Willem et al. found an increase in peer-reviewed, published papers from 38 in 2006 to 115 in 2016. Of the 698 papers covered in their analysis, only 26 included TB in some way. (39) According to Kasaie et al., “The literature on application of agent-based simulation models in the study of TB” is “quite limited”. (40)

The TB ABMs we reviewed can broadly be divided into three categories: (i) ABMs investigating the impact of interventions relating to TB diagnosis, prevention, treatment or vaccination, (ii) ABMs investigating costs or cost-effectiveness (we don't go into much detail on these since cost-effectiveness is outside the scope of this dissertation), and (iii) ABMs designed to elucidate transmission dynamics, rather than modelling interventions (also not a key focus of our work).

Much of the TB intervention modelling we reviewed was concerned with TB case-finding strategies such as contact tracing and testing people for latent TB infection. A systematic review published in 2013 found five TB models (of all types) that included contact tracing (41). Of the five models, only one was an ABM. None of the five models included sub-clinical TB as an intermediary stage between latent TB and active TB disease, as done in the TB version of ABM Spaces. There has, however, been more contact tracing TB ABMs since that review. A study published in 2014 examined the impact of contact tracing and preventive therapy in an ABM consisting of 2,000 households with one to 10 members each. Agents in this model encountered other agents in their household and select other agents in the community. The model had a timestep of one month. The model did not include other settings such as schools or workplaces and disease progression did not include a sub-clinical state. (42) Another ABM published in 2013 designed to explore the efficacy of contact tracing interventions modelled populations of 15,000 agents in an algorithm-generated network structure, that did not explicitly model households, workplaces or other environments. (43) Similar testing and contact tracing interventions were explored in an ABM closely modelled on a Canadian community of just under 9,000 people. This model contained households split into seven different communities, with some limited movement between communities. Disease progression in this model included latent and active stages, but not a sub-clinical stage. (44) Another ABM explored the impact of interventions like early diagnosis and isolation in simulated populations of just over 100,000 agents situated on a square lattice (317 by 317), with simulations run for 100 years in model time. (45) This latter study is the largest we've found in both terms of the size of the population modelled and the period of the simulation. A study published in 2019 estimated the impact of testing people in California for latent TB and treating people who test positive with TB preventive therapy. The ABM used in this study had a population of 30,000, but agents were not divided into households, workplaces or other settings. Agents in this model were more likely to have contact with agents of the same race. The risk of TB in this study was impacted by a variety of factors, such as whether an agent smokes or has diabetes. Each of the study's scenarios were run 250 times. (46)

We found several examples of ABMs being used to estimate the cost-effectiveness of TB interventions – including for treatment shortening (45-46) and urine testing in hospitals,(47-48), but relatively few that focussed on the clinical impact of novel new treatments or diagnostics. One study published in 2017 estimated the impact a hypothetical TB vaccine may have if provided to gold miners in South Africa. The study divided agents between the mine and the community and agents could move between the two settings. Agents in the model could be HIV positive, which increased their risk of developing TB. The various scenarios in the model were run for 20 years – and given this long

duration, the model also included births and deaths. A total of 216 000 model runs were used for the study. (51)

TB ABMs that model at facility or ward-level are rare. One such model estimated the impact of different interventions, such as opening doors and windows, on the risk of TB transmission at health facilities in South Africa. (52)

In some cases, TB ABMs were not used to explore the impact of interventions, but to better understand epidemiological dynamics. One study, for example, used an ABM to explore the timing of TB transmission after an agent becomes infectious and estimated that nearly half of TB transmissions occur after an individual has already been infectious for over one year. (40) ABMs were also used to estimate the timing of TB disease progression and recovery (53) and to explore the impact of multiple exposures to TB, or so-called complex contagion. (54) Finally, an innovative ABM was used to estimate the impact of poverty, poor nutrition, and poor access to healthcare services on TB outcomes. (55)

Table 1: Key characteristics of recently published TB ABMs

Study	**Research area**	**Model environment**	**Population size**	**Timestep**	**Model runs per scenario**
Tien et al. 2013	Contact tracing	Algorithm-generated network	15,000 agents	1 year (but with com-plications)	900
Kasaie et al. 2014	Contact tracing and preventive therapy	Households and community	2,000 house-holds (roughly 10,000 agents)	1 month	1,020
Tuite et al. 2017	Testing and contact tracing	Households and communities	8,952 agents	NA	100, 200, 1,500, and 5,000)
Espindola et al. 2017	Early detection and isolation	Square lattice	100,000 agents	1 day	NA
Goodell et al. 2019	Testing and preventive therapy	No structure (More likely to encounter same race and nativity)	30,000 agents	1 month	250

Caption: This table shows selected TB ABMs published since 2010 that modelled TB interventions. It excludes cost effectiveness studies and studies that used TB ABMs to explore transmission dynamics in the absence of interventions.

Our review suggests the following broad conclusions:

Firstly, comparatively few papers based on TB ABMs have been published. Those that have been published were often concerned with early case finding interventions, the cost effectiveness of interventions, or with elucidating aspects of the epidemiological dynamics of TB transmission.

Secondly, while there is some variation in the size of populations modelled in TB ABMs, population sizes were mostly in the range 2,000 to 100,000. If there are TB ABMs that model populations of hundreds of thousands or millions of agents, they are rare and did not appear in our search results.

Thirdly, we found significant variation in how TB ABMs are structured and the rules by which agents encounter each other. Model spaces included, among others, a square lattice, an algorithm-generated network, and several models that associated agents with specific environments such as households. Of the TB ABMs with explicitly defined environments such as households, models often contained only one or two environments – e.g. household and community or the mines and the community.

Fourthly, we found wide variation in the complexity of agents and the interplay of agent attributes with the risk of TB transmission or TB disease. In some TB ABMs, agents had long lists of attributes that impacted the risk of TB disease, including smoking, diabetes, and HIV status, while in others, agents had very few attributes. The risk of TB transmission was mostly determined through relatively simple equations, but in at least one instance (52) the more complex Wells-Riley formula was used.

On the whole, TB ABMs had greater complexity regarding agent attributes and transmission dynamics than we anticipated, but less complexity than anticipated regarding social structure. One reason for this may be that reliable demographic estimates can often be used to, for example, determine how many agents in a population should be diabetic, while reliable estimates of contact patterns and social structure are much harder to find and to calibrate for. In addition, complex social structures or contact patterns may also be more difficult and time-consuming to programme efficiently.

How ABMs have been used to explore COVID-19 interventions

We conducted a literature review of how ABMs are being used to explore questions relating to COVID-19 interventions. As with TB, we focus on the types of research questions that ABMs are used for, the model worlds in which agents interact, the complexity of the models, and other key design choices and trends. We searched PubMed and Google Scholar for the terms “COVID-19” or “SARS-CoV-2” and “agent-based model” or “individual-based model” or “microsimulation”. We restricted search findings to peer-reviewed articles and sifted out articles that were not about COVID-19 interventions. We paid particular attention to articles published in leading medical journals. We included systematic reviews in our analysis and draw heavily on them in the below section, given their depth and quality. Our review is not systematic in the sense that we did not count and classify all models found in our search – rather than exact quantification, our purpose was to identify broad trends in the use of ABMs for modelling COVID-19 interventions.

Since the emergence of COVID-19 in 2020, there has been a rapid proliferation of studies using ABMs to explore questions relating to the spread and containment of SARS-CoV-2. A review study by Lorig et al. published in June 2021 found that such ABMs are "widely used" and identified 126 studies (including preprints) that used ABMs to explore COVID-19 transmission processes. A review by Zhang et al. published in August 2022 identified 275 COVID-19 papers based on ABMs. (56)

Lorig et al. wrote that: "Existing models are very heterogeneous with respect to their purpose, the number of simulated individuals, and the modelled geographical region as well as how they model transmission dynamics, disease states, human behaviour, and interventions. To this end, a discrepancy can be identified between the needs of policy makers and what is implemented by the simulation models. This also includes how thoroughly the models consider and represent the real-world, e.g., in terms of factors that affect the transmission probability or how humans make decisions."(57)

The Lorig review found that 94.4% of the COVID-19 ABMs they identified modelled the impact of NPIs (at the time they conducted their literature search in late 2020 pharmaceutical interventions such as vaccines and antiviral treatments were not yet available). (57) By the time of the review by Zhang et al. there were many more COVID-19 ABMs exploring pharmaceutical interventions. (56) Lorig et al.'s review found 47.6% of studies analyse the effects of quarantining and isolation of (potentially) infected individuals and 44.4% analyse social distancing. "Testing (33.3%) and tracing (31.0%) are often analysed together, such that tracing of contacts results in them being tested and quarantined if necessary. More than half of the articles that implement either of these two NPIs (54 articles, 42.9%) include both of them," they wrote. (57)

In our review we found many studies that explored specific interventions such as the impact of closing or reopening schools (56-57) or contact tracing (60) or that explored SARS-CoV-2 transmission in specific settings such as hospitals, (59-60) supermarkets, (63) care homes (64), schools, (65) or a shop and bar (66). Other ABMs cast the net wider and modelled at the city) or national (67-68) level.

In terms of how COVID-19 ABMs handle disease states and disease progression, the Lorig review found that 63.5% of models include an incubation period, 47.6% distinguish between being symptomatic and asymptomatic (which can either be consecutive or exclusive), 34.1% define states for severely ill, and 23% for critically ill individuals that require hospital or ICU treatment.

Transmission risk is handled in a wide variety of ways and with widely varying complexity in the models surveyed by Lorig et al. They found that 27.8% of the surveyed models include location-specific transition probabilities. (57) While the processes determining whether or not transmission takes place are generally relatively simple, our review identified highly sophisticated approaches in some studies, particularly if the focus is on understanding of transmission in a specific setting.

One study, for example, used a modified version of the Wells-Riley formula to model transmission in a shop and bar. (66)

The world in which agents interact also differed widely. Unlike with TB, a substantial number of COVID-19 models are spatial models that attempt to recreate key elements of how and where people interact in the real world.

Finally, the size of modelled populations in the Lorig review ranged from as little as six to as much as 60,000,000. They found that, of the models that provide information on the population size, 7.3% consist of less than 500 agents and 12.8% of more than 5,000,000 – with the majority in between. The median population size in their sample was 30,000. “The considerable difference in population size results from a multiple of the aforementioned aspects. Usually, a trade-off needs to be made between the level of detail of the modelled individuals, as well as the complexity of their behaviour, the time step size, and the number of simulated individuals … Finding a suitable balance between level of detail and number of individuals strongly depends on the interventions that are studied.” (57)

Chapter 3 – Methodology: The ABM Spaces model, choices behind its design, and its calibration

In this chapter we provide a description of how ABM Spaces works, and we explain the reasoning behind the various design decisions and why we consider ABM Spaces to be an appropriate model with which to explore epidemiological dynamics relating to interventions for TB and COVID-19. We also provide some detail of the coding of ABM Spaces and show in detail how both the TB and COVID-19 versions of the model are calibrated to real-world data.

What is ABM Spaces?

ABM Spaces is a complex agent-based model (ABM) coded in C++. The aim of ABM Spaces is to allow for the modelling of infectious disease interventions such as contact tracing and various testing strategies in heterogeneous virtual populations that contain significant and relevant social structure. We have developed both TB and COVID-19 versions of ABM Spaces.

The term "spaces" is used to denote the fact that transmission is not directly from agent to agent, but from agent to agent via environments or spaces. This means that at any timestep in the model every environment has a transmission risk associated with it that all agents in that environment are exposed to. The degree of risk in an environment is determined by the number of infectious agents present in that environment on that timestep and by a setting-specific inherent risk factor (described later).

These spaces, or environments, are what gives the ABM Spaces' model world its social structure. There are five types of spaces in the model: households, workplaces, school classes, city blocks, and taxis (a place-holder for public transport more generally). The first three represent relatively assortative forms of mixing, while the latter two provide the model with more generalised forms of mixing. These five spaces were chosen since they have been used in real-world studies of TB transmission – our blocks space can also be classified as "other households". (71) These spaces are best thought of as groupings – that is to say that agents are associated with them, but do not move between them in a two- or three-dimensional space as would be the case in a spatial model. An agent of school-going age would, for example, be associated with a specific household and a specific school class and would be exposed to the transmission risk in those spaces.

At the start of every simulation, the model generates a series of specific households, blocks, school classes and workplaces – so that each household, for example, has a unique identifying number and a list of the agents associated with that household. These spaces are generated in such a way as to meet certain demographic parameters, so that, for example, the number of agents associated with workplaces is in line with the employment rate. Below we provide more details on the five types of spaces:

Households: All agents in the model have to belong to a household. Agents are divided into households based on a user-defined average household size parameter (the size of individual households is determined by sampling from the range 2 to (2 * average household size) - 2.). When new agents are born, they are allocated to existing households. The number of households in each simulation thus remains stable.

- Workplaces: Based on employment statistics, a certain percentage of agents are allocated to workplaces. The size of individual workplaces is determined using a similar sampling process to that used for households, although the model makes provision for much larger workplaces than households. When employed agents die or reach retirement age, they are replaced in their workplace by an unemployed agent of working age. The number of workplaces and employed agents thus remains stable in each simulation.
- School classes: All agents of school-going age are allocated to school classes. The size of school classes is determined by a user-defined parameter (set to a default maximum of 40). At the end of each year the school classes containing school-leavers are destroyed and new classes are created for agents newly starting school.
- Blocks: Blocks are simply collections of households. In the default configuration 40 households constitute one block. The utility of blocks is to allow for some limited generalised mixing of the kind one may find between people who live on the same city blocks in the real world.
- Taxis: Taxis are newly generated on every timestep so that agents can be in a taxi with different people on different days or in different weeks – thus allowing for a greater degree of generalised mixing in the model than that achieved through blocks. Taxis can be seen as standing in for public transport and similar forms of generalised mixing (the title taxis is used since taxis are the most common form of this kind of mixing in the type of community we are interested in). Only a certain percentage of the population will use taxis in any given week. It is more likely for employed persons and school kids to take taxis.

Both the TB and COVID-19 versions of ABM Spaces contain all five of these environments, but since the COVID-19 version is run for less than a year, it does not include births, natural deaths, and starting or leaving school or employment.

Agents mostly remain associated with the same spaces over time. In some instances, agents can temporarily be removed from a space – for example when an agent who is sick temporarily stops going to work or when an agent is isolated or quarantined following a positive diagnosis or by having been identified through contact tracing. More permanent removal from spaces can occur through death, leaving school, or retirement. Agents can enter spaces by starting school or entering a new job.

As noted in chapter 2, most TB ABMs that we reviewed that make use of such environments contain only one or two environments. Depending on the research question, one or two environments may be sufficient, but for certain types of research questions four or five may be needed. This may be

particularly important for interventions such as contact tracing, which in the real world is highly dependent on social networks. Indeed, in their systematic review of TB models used to evaluate the impact of contact tracing, Begun et al. advise: “As contact tracing is necessarily an activity based around individuals, more detailed approaches require inclusion of individual characteristic in models. Elements which may need to be considered include incorporation of close as well as casual contacts, location-based clustering (such as within households, schools or workplaces), age-related associations, communal contact structures, and historical contact information.” (41) While ABM Spaces does not meet all these recommended criteria, its implementation of “location-based clustering” in “households, schools and workplaces” is more complete than any other TB ABM we have reviewed.

Regarding social structure and the model world, the COVID-19 version of ABM Spaces is harder to distinguish from COVID-19 ABMs in the literature given that these vary widely and there are relatively few norms. Our impression, however, is that most COVID-19 ABMs in the literature either contain significantly less social structure than ABM Spaces, or involve significantly more complex model worlds, often with a spatial element and calibrated using geospatial data, for example from mobile phones. Though not unique, COVID-19 ABMs in which transmission is mediated through more than two or three types of spaces or groups, as in ABM Spaces, are relatively rare and might be seen as something of a middle-ground in terms of model complexity.

ABM Spaces can be adapted and parameterised to explore infectious disease dynamics for several different infectious diseases where transmission is mediated through spaces and parameterised for a wide variety of populations. In the work presented here it has been parameterised to simulate a hypothetical South African township community very loosely based on Masiphumelele outside Cape Town. Masiphumelele can be seen as a stand-in for a generic, relatively closed South African township community of 10,000 to 20,000 people and our work should thus have some generalisability.

Since ABM Spaces contains substantial stochasticity, we report means and uncertainty ranges for all key findings. Unless otherwise stated, all ranges are 95% ranges (percentile 2.5 to 97.5). For the analyses in chapters 4 and 6 each scenario is run 10,000 times. For the analysis in chapter 5 each scenario is run 1,000 times (see chapter 5 for the reasoning behind this lower number).

Model structure

Though the structure of ABM Spaces is somewhat complex, its basic building blocks are straight-forward.

As described above, the model world consists of a large number of spaces, each with its associated agents. Though this model world can and does change somewhat over time (more so in the TB version of the model), it is largely determined at the beginning of each simulation.

In this model world two types of processes, pathogen transmission and disease progression, then begin and continue on each timestep for the duration of the simulation. The pathogen transmission process involves counting the infectious agents in each space in the model and then exposing uninfected agents in those spaces to the appropriate risk of becoming infected. The disease progression process involves letting disease progress appropriately within each agent, for example letting an agent's latent TB infection progress to sub-clinical disease after a certain number of weeks. In this way, an agent may for example as part of the pathogen transmission process become infected with TB due to a high risk in his school class. The disease progression process within that agent will then be started and the agent may accordingly become infectious several months later. Once infectious, the agent will then contribute to the transmission risk in the settings with which he is associated. In this way, the pathogen transmission and disease progression processes together, and in the interaction between them, form the basis for our model of the epidemiology of an infectious disease in a community.

Attributes of agents in ABM Spaces

Each agent in ABM Spaces has a lengthy list of attributes. These include the agent's age, the agent's association with specific environments, and a series of disease-related attributes or state variables. Agents in the version of ABM Spaces adapted for COVID-19 (described in chapter 6) have a set of COVID-19-related attributes – for example defining whether or not they are infectious on a specific timestep. Agents in the version of ABM Spaces adapted for TB (described in chapter 4) have a longer set of attributes relating to TB and HIV disease states and susceptibility. HIV is included here, since uncontrolled HIV infection increases the risk that someone with latent TB infection will develop active TB disease. Neither version of the model currently includes gender, though the meaningful inclusion of gender is a priority for future work.

In the TB version of ABM Spaces, the model lets agents age over time. This is because TB disease progression is very slow and the meaningful modelling of interventions often require the simulations of multiple years (10 years in the TB model). Since COVID-19 is a much faster moving infection, the COVID-19 version of the model only runs for 270 days and accordingly agents in that model do not age. Leaving the ageing routine out of the COVID-19 version of the model is an example of how a model can be simplified, and speeded up, to suit the question being explored.

As explained in the previous section, agent associations with environments can change over time, either due to ageing, death, quarantine and isolation, and, in the COVID-19 version, hospitalisation.

Finally, ABM Spaces also keeps track of a lengthy list of indicators relating to each agent. These are mostly to allow for tracking and analysis. For example, the model keeps track of the type of environment in which each agent with TB infection became infected, thus allowing for aggregate

statistics on where agents got infected – which, as we will see, is important for model calibration. The model also keeps track of things like how many times an agent has been tested.

Population size

ABM Spaces can be configured to simulate populations of any size, but in the TB and COVID-19 analyses populations of 20,000 and 10,000 are used. Our choice of population size is primarily driven by the fact that our aim is to model the impact of interventions at the level of a single community or township. This is in part because (i) calibrating the model at the level of a single community can be done more credibly than for larger more complex communities, but also because (ii) in our judgement a population of 20,000 is sufficiently large for the meaningful exploration of the dynamics we are interested in (additional agents would add complexity and slow the model down without providing benefit) and (iii) we are sceptical that the relative uniformity of a population of 20,000 or 10,000 could be assumed for larger populations – which is to say that when creating larger populations simply increasing the number of agents would not be sufficient and to keep the model roughly realistic various sub-populations with, for example, variable employment rates, may be required (once again introducing additional complexity without clear benefit for exploring our research questions).

As noted in chapter 2, most TB ABMs we reviewed had populations in the range 2,000 to 100,000, while populations in COVID-19 ABMs ranged from 6 to 60,000,000 (with a median of 30,000). Our TB and COVID-19 analyses thus falls squarely in the typical range of population sizes in TB ABMs.

Given the population sizes used in our work there is a possibility that stochastic fade-out may occur in some model runs. We have however not observed this happening, likely due to the fact that the number of initial infections in our model is high enough to make stochastic fade-out extremely unlikely. We have included uncertainty ranges for all key model outputs, in part to provide insight on the presence and absence of such unexpected effects.

Modelling two processes

As argued in chapter 2 and briefly described in the ‘model structure’ section above, the two key processes in this type of ABM are the process by which disease progresses in individual agents and the process by which pathogens are transmitted from agent to agent within a structured model world. Modelling these processes simultaneously, efficiently, and in an interconnected way is the key programming challenge and source of complexity in ABM Spaces.

The disease progression process in ABM Spaces works as follows: On the timestep that an agent becomes infected with a disease pathogen a series of disease progression parameters are drawn from probability distributions based on the published literature. The parameters indicate whether and

on which dates certain disease progression events, such as becoming infectious or requiring hospitalisation, will occur in the absence of treatment (in chapters 4 and 6 we provide details of how these processes work in the TB and COVID-19 versions of the model). In this way, an agent's future disease trajectory is mapped out on the timestep that the agent becomes infected. These dates are not set in stone and can be changed by interventions – so that, for example, an agent that was scheduled to die may live if diagnosed and started on treatment. There are two advantages to mapping out an agent's future trajectory like this – as opposed to an alternative whereby no future dates are drawn and agents are simply exposed to certain risks on every timestep. Firstly, it is computationally much less intensive since sampling happens once rather than on every timestep. Secondly, it allows for the relatively straight-forward sampling from appropriate probability distributions – though essentially the same effect can be achieved with sampling on every timestep, doing it in that way would require the extra step of transforming the probability distributions into daily risks.

In some important respects, the TB disease process in ABM Spaces is more detailed than TB ABMs reviewed in chapter 2. For example, none of the TB ABMs that we reviewed included a sub-clinical TB stage as we do in the TB version of ABM Spaces. Also, while some TB ABMs that we reviewed did allow for sampling from probability distributions, this was by no means the norm. By contrast, disease progression in the COVID-19 version of ABM Spaces is fairly typical, although including an asymptomatic state and hospitalisation places it toward the more complex/realistic end of the spectrum.

As mentioned earlier in this chapter, the transmission processes in ABM Spaces are not from agent to agent, but from agent to space to agent – or more precisely, infectious agents increase the risk in specific spaces and other agents associated with those spaces are then exposed to that risk. On every timestep, the model calculates the transmission risk for every specific space. The transmission risk is the product of two factors: the number of infectious agents associated with that space, and the inherent risk parameter associated with the type of space. The five types of spaces each have different inherent risk parameters (the calibration section below explains how these inherent risk parameters are calibrated). On each timestep, the model cycles through all spaces and exposes all uninfected agents to the transmission risk in those spaces. Should transmission take place, the newly infected agent is started on the disease process (described above) – which in time may lead to the newly infected agent also becoming infectious and in turn exposing members of its associated spaces to an increased transmission risk.

To determine whether or not an agent becomes infected on a specific timestep in a specific space, the model draws a random number between 0 and 1 and then checks whether the number falls within a range that would let infection occur (in line with the approach described by Vynnycky and White). (1)

In both the TB and COVID-19 versions of the model, the range is determined by the following formula:

$$r = 1 - (1 - i)^n$$

Where *n* is the number of infectious agents in that space on that timestep and *i* is the inherent risk parameter for the specific type of space, which is always a number between 0 and 1.

There is thus a value of I for each type of setting in the model – i_h for households, i_w for workplaces, i_s for schools, i_t for taxis, and i_b for city blocks. If the inherent risk parameter for households is 0.05 and there are two infectious agents in a specific household on a specific timestep, the risk range in that household on that timestep would be calculated as follows: $r = 1 - (1 - 0.05)^2 = 0.9025$. For each member of the household a random number between 0 and 1 would be drawn and if that number is lower than 0.0975 the agent will become infected.

As explained in more detail in the calibration section, the five values of I are not directly parameterised based on published data. Instead, the five values of I are simply those values that result in the right number of infections occurring over time in the various types of settings, as estimated in the literature.

The formula is slightly different for some settings and in some cases agents can be treated differently. In the TB version, for example, agents with sub-clinical TB are half as infectious as agents who are symptomatic. Such agents contribute only 0.5 to the n in the $r = 1 - (1 - i)^n$ calculation. ABM Spaces defaults to the Reed-Frost equation since it is recommended by Vynnycky and White (1) as the standard for such models, but it could easily be adapted to apply novel formulas.

In the COVID-19 version of ABM Spaces the risk of dying from COVID-19 is impacted by an agent's age. In neither the TB or COVID-19 versions however is transmission risk directly impacted by age, though there is some indirect impact since the calibration of the setting specific risk is based on estimates of real-world transmission of TB in such settings, which will in part be driven by differing levels of infectiousness between for example school-going children and working adults. Current versions of ABM Spaces include gender, but gender does not impact transmission risk or disease progression. Adding this is a priority for future work on the model.

Rationale behind network properties, heterogeneity and stochasticity in ABM Spaces

As discussed in chapter 2, the worlds within which agents in ABMs interact can take many shapes and be constructed in many ways. Our choice of a model world defined by spaces is firstly determined by the fact that, on the face of it, such a design corresponds well to the real world. People after all live in households situated in city blocks or neighbourhoods, and children go to school and some adults work. That we have good data on the likelihood of TB transmission taking place in these spaces strengthens the case. We'd argue that such a design more accurately corresponds to social networks in the real world than, for example, an ABM in which agents are situated on a line or square lattice where they interact with neighbouring agents. It might be argued that a two- or three-dimensional spatial model world in which agents move around could correspond even more closely with reality, but

such a model would be harder to calibrate reliably and we are sceptical as to whether the extra model complexity would provide additional insight into our research questions.

A design such as that used in ABM Spaces gives rise to more complex social networks than may at first be obvious. Consider that a single household of six agents may contain two agents that are also associated with school classes and one agent that is also associated with a workplace – all six agents will be associated with the local block and some may from time to time take public transport. There are thus multiple roots by which infection can be introduced into the household and exported from the household. Each such household, workplace, or other space in the model can be thought of as a node in a network, with connections between two nodes existing where an agent is associated with both – for example, an agent that is associated with household 3 and workplace 14 would connect those two nodes. In addition, while most nodes are relatively stable over time, taxis are constantly reconstituted and are thus transient nodes. Since blocks are collections of households, they are a type of super node in terms of size, although the very low infection risk associated with blocks means that blocks have only very weak links with other nodes. In simulations of 20,000 agents, ABM Spaces thus produces complex networks with hundreds of nodes of various sizes and many thousands of connections of various strengths between nodes.

One strength of ABM Spaces is that the constituent parts of these complex social networks, the individual agents and the association with specific spaces, are newly created every time the model is run. You may thus have some runs of the model with significant clustering (where a few households are highly connected and others are poorly connected), and some runs with less clustering (where households are more equally connected). These variations in network structure likely reflect the kind of random diversity in real-world populations.

The variation in network structure between runs of the model is one source of stochasticity, or randomness, in ABM Spaces. Another stochastic element is the fact that the initial infections in every run of the model are randomly distributed – that is to say the model randomly picks which agents to infect when the model is initialised. It is thus possible that in some runs of the model a disproportionately high number of the initially infected agents will happen to be in highly connected households or workplaces, which would then be likely to result in more rapid, and possibly larger, outbreaks. Similarly, if initial cases are disproportionately situated in poorly connected households or school classes, the resulting outbreak may be smaller or take longer to gather momentum. Either way, the random distribution of initial infections combined with the generation of new networks on every model run produces significant stochasticity.

These complex networks provide a useful model for studying the spread of disease within social networks, but as importantly, they also provide a model within which to explore the efficacy of interventions such as contact tracing or early testing. Thus, in ABM Spaces contact tracing is not simply a matter of adding an extra parameter to an ordinary differential equation, it involves tracing the specific agents who share a household or school class with a newly diagnosed agent and then testing or quarantining those agents. The quarantining and isolation of agents thus occurs within

these networks and can in some cases have the effect of temporarily shutting down the links between specific nodes.

Several other processes in the model are also stochastic in nature. As described above, whether or not transmission takes place in a specific instance is determined by checking whether a random number falls within a certain range. It is thus possible that, purely by chance, an outbreak in a certain space might be slowed if a series of draws happened to fall outside of the range required for transmission. The disease process in individual agents is also highly stochastic, since for each agent a series of dates are sampled from probability distributions and these dates determine the disease course. Here to, while generally disease course will tend to follow the distribution described in the literature, it is possible that, purely by chance, an unusually high number of agents may stay infectious for longer and thus contribute to an outbreak being larger than it might have been otherwise.

The combined impact of these various stochastic processes is that there can be significant variation between different runs of the model. Accordingly, running the model only a few times is likely to produce non-representative outputs purely by chance. This is why in all the core analyses, scenarios are run 10,000 times and mostly mean output values are used.

Model calibration and parameterisation

The calibration and parameterisation of ABM Spaces can be divided into three broad categories.

Firstly, the model is parameterised so that the model population meets certain demographic indicators. This includes parameters such as employment rate, average household size, maximum size of school classes, age distribution, and in the TB version, HIV status. A number of reliable sources exist in South Africa for these demographic indicators and they are referenced in chapters 4 and 6. National level indicators have mostly been used given the lack of accurate community-level figures.

Secondly, disease processes are parameterised based on the published literature. Here the model samples from probability distributions rather than using fixed parameters. In the TB version for example, discreet probability distributions were constructed based on the literature describing time from infection to active disease. In this case, the distribution was such that most agents would progress to active disease in a matter of months, but a small minority of agents would progress only years after infection. As with the demographic parameters, the aim was to make processes in the model world correspond with the reality of disease processes (as described in the literature) in a number of key areas that we consider relevant to the epidemiological dynamics being modelled. While we only mention it briefly here, the careful parameterisation of disease processes based on the published literature is an essential and time-consuming part of the methodology behind ABM Spaces – without it, the model would be an empty shell. In the more detailed calibration sections in chapter 4

and 6 we provide a detailed description of how the TB and COVID-19 versions of ABM Spaces were parameterised and calibrated based on the published literature.

Thirdly, ABM Spaces uses a somewhat novel approach for the calibration of transmission risk. Rather than hard-coding a value for the five setting-specific values of I in the transmission formula $r = 1 - (1 - i)^n$ the five values of I were varied so that the model hit certain targets regarding how many infections take place and where.

For these targets we firstly drew on published research estimating what percentage of TB infections in a Cape Town community took place in households, schools, workplaces, public transport, and in the community. By integrating data on social interactions and environmental context, Andrews et al. developed a novel approach to project the impact of age-specific contact patterns and estimate where TB transmission occurs. They utilized local data from a study of social interactions together with measurements of carbon dioxide in common indoor environments, to model TB transmission in a South African township. (71) While Andrews et al. reported what percentage of transmission occurred in which of the five settings by age group, they did not provide overall numbers (irrespective of age) for all five of the types of settings. Although the authors could not provide these overall percentages on request, the numbers can be calculated based on the data that was published. Accordingly, we calculated that 15.6% of transmission occurred in households, 41.4% in workplaces, 10.2% in schools, 21.9% in public transport, and 10.9% in the community.

We also drew on estimates of disease spread – WHO estimates of TB incidence in the TB version and early published estimates of R_0 in the COVID-19 version.

Our six target values were thus the five percentages and the measure of spread (TB incidence and the R_0 value of COVID-19, respectively).

In order to calibrate the model to produce outputs as close as possible to these six target values, we varied the five setting specific risk values (i_h , i_w , i_s , i_t , and i_b) used in the transmission risk formula $r = 1 - (1 - i)^n$. The inherent risk parameters *i* for the five types of spaces are thus simply the parameter values required to calibrate the model so that the right number of infections take place in the right places.

There are a wide variety of methods used for the calibration of ADMs,(72) with the choice of method often determined by factors such as model size and complexity. Options range from manual parameter adjustment, to mathematical methods such as maximum likelihood estimation, to several more complex approaches used in economics ABMs. (73) Both the TB and COVID-19 versions of ABM Spaces were calibrated by manually varying the values of i_h , i_w , i_s , i_t , and i_b). We opted for manual calibration firstly because it is safe to do so given the specifics of our model. The key risk with manual calibration is that there may be multiple solutions and that calibration might by chance limit the model to one non-representative solution. Given the nature of the variables and the target values we judged the risk of this occurring with ABM Spaces to be very low. The five values of I relatively directly (via the transmission formula) influences the percentage of infections that take place in each

of the five types of settings. The fact that the five percentages have to be hit simultaneously with TB incidence or COVID-19's R_0 value, acts as a severe constraint on potential solutions for the five I values. In both models, our data show that our modelling only covers one, or part of one wave – the only real variation is thus whether the same sized epidemic can be achieved with a longer and flatter versus a shorter and steeper epidemic. While this may in part be influenced by values of I, it also depends on a series of other parameters such as how soon after exposure agents become infectious and how long they remain infectious (both parameters that are directly drawn from the published literature). Secondly, while it is possible to create an automated calibration routine, executing such a routine is likely to be time-consuming. Since both versions of the model have significant stochasticity, hundreds of runs of the model are needed to get a clear picture of the impact of a set of values for I (in all the key analyses in chapters 4, 5, and 6 each scenario is run 10,000 times). Even though ABM Spaces is fast, automated calibration routines would likely require a few hundred thousand runs to complete on a standard desktop computer and would thus be much less efficient than the manual process we employed. (In the interest of reproducibility, the code used for the calibration process is included in the GitHub repository.)

In chapters 4 and 6 we provide more detail on the setup and calibration of the TB and COVID-19 versions of the model.

Coding of ABM Spaces

While ABM Spaces was prototyped in R, the final versions of both the TB and COVID-19 versions of ABM Spaces were developed in C++. C++ was chosen mainly to allow for maximum speed – the R versions were unacceptably slow. Since these models are highly complex, contain significant heterogeneity and have to be run hundreds of times to get meaningful results, we did not consider slower or higher-level programming languages or packages to be an option since slower speeds would have severely limited the utility of the models and would likely have forced compromises in terms of model complexity. The final C++ versions of these two models allow hundreds of simulations to be run in a few minutes on standard hardware.

Much of the code in the TB and COVID-19 versions of ABM Spaces is the same. Both versions, for example, use the same code to generate households, workplaces and school classes and to associate agents with these spaces. There are, however, also substantial differences between the two versions. The parameterisation of disease processes and the number and logical flow between disease states are different. In addition, while hospitalisation was highly relevant for the COVID-19 model, it was not considered essential for TB and left out of that version. On the other hand, since the TB version has to run for 10 years, it contains routines for ageing, leaving or starting school, and retirement and employment, which are not in the COVID-19 version, which runs for less than a year.

The TB version also contains extra code to implement X-ray TB screening and targeted universal TB testing. Finally, there are also differences in the outputs produced by the two versions.

At the core of both models is an array of agents, each with a long list of attributes and relevant data fields tracking various changeable states. One key to the speed of the model is that most operations on the array of agents happen in place (without copying). The model is designed to run in parallel – thus allowing for higher speeds on multicore systems.

The basic flow of the programme is as follows:

Create agents

Create spaces and allocate agents to spaces

Start loop of timesteps

Let disease progress in individual infected agents

Calculate transmission risk in each space and expose associated agents to that risk

Run interventions (depending on scenario)

Write interim output data

End loop of timesteps

A number of command line switches are available that allow users to pass arguments specifying for example how large the population should be (num_agents), how many days (COVID-19) or weeks (TB) to model (num_time), and how many cores to use (cores). A full list of switches is included in the source files. There is also a Makefile and a number of R scripts that can be used to run the models and generate output in HTML format using the RMarkdown package. Users who have cloned the ABM Spaces repository can for example run “make analyze 1000” in a Linux terminal to reproduce the key analysis described in chapter 4 (it requires R and the Rmarkdown, dplyr, and HTMLTable R packages to be installed). The source code, scripts, and model outputs used for the analysis presented in chapters 4, 5 and 6 are all available at https://github.com/marcuslowx/ABM-Spaces.

Chapter 4 – Evaluation of four tuberculosis case-finding interventions in a complex agent-based model

In this chapter we present an implementation of ABM Spaces used to explore the impact of four TB case-finding interventions on tuberculosis (TB) detection rates, incidence and mortality. The analysis presented illustrates our core contention that relatively complex agent-based models can meaningfully be used to explore real-world epidemiological questions and that by coding in a compiled, low-level language such as C++, such models can be made very fast.

Background

As noted in chapter 1, people with TB are estimated to transmit the bacteria to three to 10 people per year if not diagnosed. Accordingly, a key question in TB control is how we should go about diagnosing more people with TB more quickly. The problem has become even more complex in recent years with evidence showing that TB transmission can occur while someone is still in a sub-clinical (asymptomatic) phase. Since the importance of sub-clinical TB has only become apparent in recent years, much of the historic literature and modelling efforts does not take its impact into account.

Several early diagnoses, or active case finding approaches, are currently being used in health systems, being piloted, or are being researched in large trials. Most notable among these are contact tracing, digital X-ray screening, and targeted universal testing (TUT). Which of these approaches, or combination of approaches, will be most effective in particular circumstances is unknown. In addition, while real-world studies can give some indication of improvements in case detection, due to the complexity and long timelines involved, it is hard to imagine studies that will be able to prove a clear link between case finding approaches and outcomes such as decreases in TB cases and decreases in TB mortality. Modelling offers one means to explore these impacts in the context of our current understanding of sub-clinical TB.

We use ABM Spaces to explore the impact of four different active case finding approaches on TB diagnosis, TB cases, and TB mortality. The four approaches are: TUT, contact tracing with X-ray screening of contacts, TUT plus contact tracing with X-ray screening of contacts (called the ultra scenario), and mass X-ray screening (a base scenario is also included for comparison).

Since there is significant uncertainty in the TB literature regarding time from TB exposure to developing TB disease and time from starting TB treatment to becoming non-infectious, secondary analyses are presented showing what impact alternative parameterisation in these areas has on model output. A secondary analysis is also presented showing the impact of higher HIV prevalence on the key model outputs.

Setup and calibration of the TB version of ABM Spaces

Below we explain the design of ABM Spaces TB in more detail, with reference to how we define agents, how disease progression works for individual agents, how we define spaces (or settings), how TB transmission works in the model, the inclusion of HIV co-infection in the model, how we model interventions, and model calibration and choice of parameters.

a. Agents in the TB model

In its default configuration, the model generates 20,000 unique agents at the start of every simulation. Each agent has a set of attributes, such as age, gender, TB status, and HIV status. Some of these attributes can change over time. Agents, for example, get older or can become infected with TB. The initial ages of agents are determined by sampling from a distribution that reflects the reported age distribution in South Africa. (74) New agents are born during simulations at a rate determined by a user-defined fertility rate. Agents die either due to TB (according to processes described lower down), or naturally, according to a user-defined death rate. The inclusion of births and deaths from causes other than TB in the TB version of the model is necessary given that simulations model at least 10 years.

b. Disease progression in the TB model

As explained in chapter 3, the two key processes we are interested in is the way disease progresses in individual agents and the way TB is transmitted from agent to agent within a model world with realistic social structure. The disease progression process works as follows: On the timestep that an agent becomes infected with latent TB, a series of disease progression parameters are drawn from distributions based on the published literature (see details in section g below). The parameters indicate whether and on which dates certain disease progression events will occur in the absence of treatment. It is determined (i) whether an agent with latent TB infection will progress to sub-clinical disease, (ii) if yes, when progression to sub-clinical disease will occur, (iii) whether an agent will develop active TB disease, (iv) if yes, when progression to active TB disease will occur, (v) whether an agent will die, (vi) if yes, when an agent will die, (vii) whether an agent will be cured in the absence of treatment, and (viii) if yes, when cure without treatment will occur. The flowchart in figure 1 sets out the paths that an agent's disease course can follow.

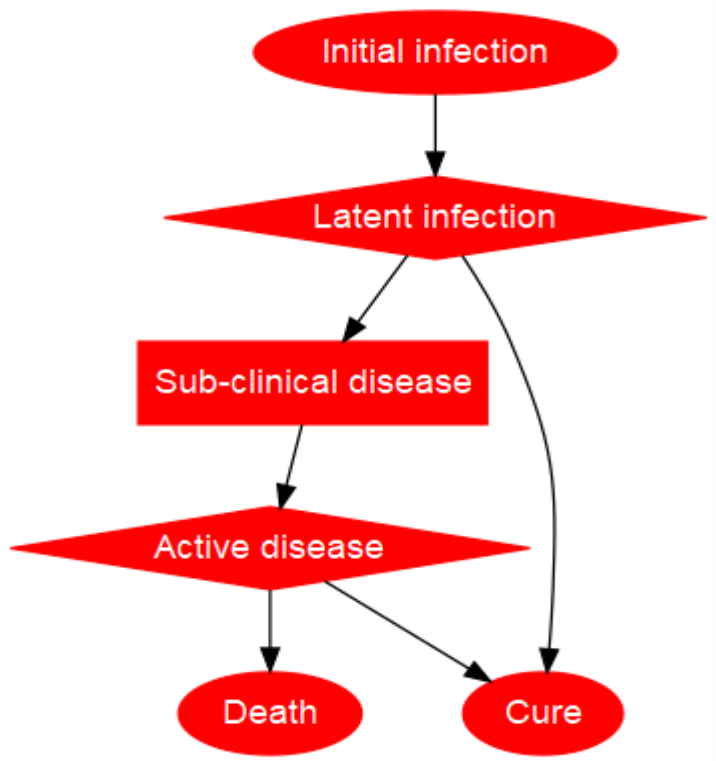

Figure 1: Flowchart of TB disease progression in ABM Spaces

In this way, an agent's future disease trajectory is mapped out on the timestep that the agent becomes latently infected. These dates are, however, not set in stone and can be changed by interventions – so that, for example, an agent that was scheduled to die may live if diagnosed and started on treatment.

c. Spaces in the TB model

ABM Spaces contains five types of spaces: households, workplaces, school classes, city blocks, and taxis (a place-holder for public transport more generally). The first three represent relatively assortative forms of mixing, while the latter two provide the model with more generalised forms of mixing. These five spaces were chosen since they have been used in real-world studies of TB transmission – our blocks space can also be classified as "other households". (71) See chapter 3 for a more detailed explanation of how these spaces are constituted and the type of network that results from their creation.

d. TB transmission

TB transmission in the model is not from agent to agent, but from agent to space to agent – or more precisely, infectious agents increase the risk in specific spaces and other agents associated with

those spaces are then exposed to that risk. On every timestep, the model calculates the transmission risk for every specific space. The transmission risk is the product of two factors: the number of infectious agents associated with that space, and the inherent risk parameter associated with the type of space. The five types of spaces each have different inherent risk parameters (section g below explains how these inherent risk parameters are calibrated in the TB version of the model). On each timestep, the model cycles through all spaces and exposes all uninfected agents to the transmission risk in those spaces. Should transmission take place, the newly infected agent is started on the disease process described in section b above – which in time may lead to the newly infected agent also becoming infectious and in turn exposing members of its associated spaces to an increased transmission risk.

Agents in the TB version of ABM Spaces with sub-clinical TB are half as infectious as agents with active TB. We could not find reliable estimates of the difference in infectiousness between sub-clinical and active TB in the literature – although it seems very likely that people with sub-clinical disease will be significantly less infectious given less frequent coughing and lower bacterial load.

e. HIV co-infection

Given that living with HIV has been estimated to increase a person's risk of developing active TB disease by more than 20 times (5) and that around 13% of the South African population are estimated to be living with HIV, we included HIV status in ABM Spaces TB. At the outset of each simulation 13% of agents are randomly defined to be HIV positive. This 13% is further divided into the following HIV stages (with the likelihood of falling into the different stages determined by outputs from version 4.4 of the Thembisa mathematical model of HIV in South Africa (3): (0) HIV-negative, (1) Diagnosed HIV-positive and virally suppressed, (2) Diagnosed HIV-positive and not suppressed, (3) HIV-positive, but not diagnosed and not suppressed. In the model, agents at stage 2 or 3 are at 20 times higher risk of proceeding from latent TB infection to sub-clinical disease than HIV-negative agents. Including information on whether agents have diagnosed HIV also allows the model to let these agents be tested more regularly in scenarios with the TUT intervention.

As time passes in each simulation, additional agents acquire HIV in line with HIV incidence estimates drawn from the Thembisa model and agents move between stages. The model thus does not directly model actual HIV transmission in the way that it models TB transmission. This is for two reasons: Firstly, the relationship between HIV and TB is one directional, with HIV status impacting TB risk, but not the other way around. The key dynamic for our modelling is simply that a certain percentage of the population should be exposed to this increased TB risk. That dynamic can be sufficiently captured without direct modelling of HIV transmission. Secondly, directly modelling HIV transmission in ABM Spaces TB will dramatically increase model complexity since it will require the current setting-mediated TB transmission framework to be augmented with a second type of direct agent-to-agent transmission framework that includes complexities such as partner matching along the lines of previous work by Geffen and Scholz. (75) Combining these two types of agent-based modelling into

one model will at best offer marginal benefits in answering the questions we are interested in here. The marginal benefit is outweighed by the increased cost in terms of complexity and model speed. Such a unified model may however be worth developing in future in order to explore other questions relating to HIV and TB co-infection.

f. Interventions in the TB model

The TB model was designed so as to allow for the modelling of the following intervention scenarios:

1. Base: 80% of symptomatic agents are tested using molecular testing, with testing typically happening in the weeks after an agent first becomes symptomatic. In the remaining 20% of cases symptomatic agents never present to facilities. There is no testing of asymptomatic agents. 80% of agents who test positive are treated and cured (this rate can be seen as the product of the linkage to care rate and the treatment success rate). This scenario can be seen as the standard of care, or control scenario. Even though in reality there is some tracing of household contacts, our understanding is that it is far from universal.
2. Base + CTI: 80% of symptomatic agents are tested using molecular testing as in the base scenario. Contacts of agents who test positive are traced and tested, with the process being repeated for any contacts that test positive. The model does not include preventive therapy and contacts who test negative are accordingly not started on preventive therapy.
3. X-ray mass screening: 80% of symptomatic agents are tested using molecular testing (as described above) and mass X-ray screening is provided so that every city block is screened once in two years. When a city block is screened 60% of agents in that city block is X-rayed. X-rays are given a 90% sensitivity for the detection of sub-clinical or active TB. This is an attempt to model a very ambitious X-ray screening campaign. It is inspired by findings from South Africa's First National TB Prevalence Survey indicating that X-ray screening can help detect many sub-clinical cases that would be missed otherwise.
4. TUT: 80% of symptomatic agents are tested using molecular testing (as described above) and 80% of high-risk agents are tested once a year using molecular testing (provided they have not already been tested as contacts of agents with TB or because they have symptoms themselves). High risk agents include agents with diagnosed HIV, agents who have had TB in the last two years, and household contacts of agents with diagnosed TB. This scenario is an implementation of the TUT strategy that has been found in a cluster randomised trial to increase TB case detection by 17% compared to standard of care.
5. Ultra: 80% of symptomatic agents are tested using molecular testing (as above), 80% of high-risk agents are tested once a year using molecular testing (as per TUT scenario above), and 80% of household, school class and workplace contacts are screened using X-rays.

In order to avoid agents being exposed to too many X-rays or an unreasonably high number of molecular tests, all the above scenarios are designed such that an agent cannot receive more than one molecular test in six months and not more than one X-ray per year. In practice this means that agents will at times sit-out screening or testing that is done as part of tracing or TUT.

g. Calibration and parameter values

Disease progress parameter distributions

As explained in section b above and in chapter 2, one of the key processes modelled in ABM Spaces TB is the TB disease progress within individual agents. Instead of letting all agents experience the same disease course, key questions about each agent's disease course are decided by drawing from probability distributions. For example, when deciding for how long an agent will be sub-clinical before developing active TB disease, the base configuration of the model samples from a uniform distribution of four to 48 weeks. The distributions used in our model are shown in table 1.

In what follows we explain the basis for our choice of probability distributions. We draw largely from a tuberculosis primer published in *Nature* in October 2016. (5) We also consulted the original sources referenced in the *Nature* primer and searched the literature for relevant papers published in the years since the publication of the *Nature* primer. The probability distributions used in the model are shown in table 2.

One challenge in modelling TB is that significant uncertainty remains concerning some of these probability distributions. It is for example uncertain what percentage of people, if any, who develop sub-clinical disease never progress to symptomatic disease. (12) We assumed that in the absence of treatment all do progress. There are also differing views in the literature regarding the likelihood that latent infection can progress to symptomatic disease two years after initial infection – one recent study argued that this is extremely rare, (76) while the more traditional view was that it is not that rare. While such uncertainties undoubtedly present a challenge when modelling TB, there is arguably a slightly higher tolerance for uncertainty when the objective of the model is to explore epidemiological dynamics associated with interventions rather than make precise predictions about potential future outcomes. Our focus here is on the prior.

We have mostly used uniform probability distributions. This is because in most cases, as mentioned above, the real values are highly uncertain and there is little evidence to support the use of more nuanced distributions. We have, however, in two instances opted to use more complex probability distributions, or provided an option to use alternative probability distributions, where using uniform distributions may fail to capture a key aspect of the real-world dynamics of TB and there is good evidence to support the use of more nuanced distributions.

1. Likelihood of progression from latent infection to sub-clinical disease: According to the *Nature* primer, 5% to 15% of people with latent TB infection will progress to active disease (5) – they cite a study estimating different levels of risk for different age groups. (77) Similarly reliable estimates are not available for progression from latent infection to sub-clinical disease. In the base configuration of our model, we have however assumed that sub-clinical infection is a state everyone passes through when progressing from latent to active TB and that everyone who develops sub-clinical TB will go on to develop active disease, providing they are not treated. For the purposes of our model, the best estimate for progression from latent to sub-clinical TB is thus equal to the best estimate for progression from latent to active disease. Accordingly, our model uses the mid-point of the *Nature* primer estimate to give each latently infected agent a 10% risk of progressing to sub-clinical TB.
2. Timing of progression from latent infection to sub-clinical disease: Since there are few diagnostic tools with which to diagnose sub-clinical TB and since sub-clinical TB is comparatively poorly studied, we have not been able to find reliable estimates of the time from latent infection to the development of sub-clinical TB. One review of insipient and sub-clinical TB provided no concrete numbers but indicated that the timing can be widely variable. It stated that, “the time from initial infection to the development of incipient or subclinical TB may range from weeks to a lifetime”. (10) Since sub-clinical TB is typically a precursor to active TB, we can, however, infer that the period from latent infection to sub-clinical TB will be shorter than that from latent infection to active TB, which has been much better studied. While we discuss it in more detail in point d below, suffice here to note that in the vast majority of cases active TB is thought to develop within a few months to two years from latent infection. Within these parameters, we make an educated guess that sub-clinical TB can develop in anything from one month to 18 months after initial latent infection. In the standard configuration of our model, we transform this to weeks – producing a uniform distribution of four to 72 weeks. As a secondary analysis, we also ran the model with an alternative non-uniform distribution that allows for active TB to develop up to a decade after exposure – although likelihood of developing sub-clinical TB in this distribution is still much higher during the first 18 months.
3. Likelihood of progression from sub-clinical to symptomatic disease: It is not known what percentage of people with sub-clinical TB will progress to active TB disease. Though there is some suggestion that a small percentage of people with sub-clinical TB will not go on to develop active TB, we suspect that this is highly unusual given the emerging picture of sub-clinical TB as a precursor to active TB. Since the number of exceptions are probably negligible, we’ve assumed in all our base scenarios that 100% of people who develop sub-clinical TB will go on to develop active TB. The model does not allow TB status to oscillate between sub-clinical TB and active TB.
4. Timing of progression from sub-clinical to symptomatic disease: According to the *Nature* primer, progression from latent infection to active TB disease can occur in anything from

months to a few years. While there seems to be relative consensus that most people who develop active TB disease do so within a few years of becoming latently infected, there is disagreement regarding the likelihood of developing active TB more than two years after exposure. A landmark study published in the *British Medical Journal* in 2018 argued that the incubation period of TB is shorter than previously thought and that the vast majority of people develop active TB within two years and that only very few do so after that. (76) This contrasts with the more traditional interpretation that people can develop active TB many years, even decades after exposure. To account for these differing views, our model includes two different distributions from which to sample when determining when an agent will develop active disease. The first (and default), uses a uniform distribution of four to 48 weeks – corresponding, though crudely, to the analysis in the *BMJ* article. The second distribution also sees most agents developing active TB in the first few years after exposure, but has a long tail, stretching to 10 years. This latter distribution more closely resembles the traditional view. (The two distributions are visualised in figure 2A.) Note that the studies cited here relate to the time from latent to active disease – since our model includes the intermediate step of sub-clinical disease, the sub-clinical period is essentially subtracted from these estimates when calculating model parameters.

5. Likelihood of progression from symptomatic disease to death: In a widely cited review by Tiemersma et al*.* it is estimated that life-time case fatality for untreated smear-positive TB is 70%, while for smear-negative TB it is over 20%. (78) The *Nature* primer states that "in the absence of treatment, approximately 50% of individuals who develop active TB disease will succumb to it". Since our model does not distinguish between smear-positive and smear-negative, and since the *Nature* primer estimate is close enough to the mid-point of the smear-positive and smear-negative estimates, we use the 50% figure in our model.
6. Timing of progression from symptomatic disease to death: According to the review by Tiemersma et al., the duration of TB from infection to cure or death is approximately three years – although for the purposes of our model parameters the latent and sub-clinical periods have to be subtracted from this number. According to the *Nature* primer the average duration of infectiousness among those with active TB disease is greater than one year in many high-burden settings. In line with these estimates, our model samples from a distribution of 24 to 72 weeks – uniformly distributed.
7. Likelihood of progression from symptomatic disease to cure (without treatment): This parameter is necessarily the inverse of that described in point e above. Accordingly, the model gives agents with active TB disease a 50% chance of being cured in the absence of treatment.
8. Timing of progression from symptomatic disease to cure (without treatment): The review by Tiemersma et al*.* states that the time from disease onset to death is roughly the same as that from disease onset to cure. We thus use the same distribution here as described in point f above. It is also in line with the *Nature* primer's statement that people with active TB

are on average infectious for more than a year – the 24-to-72-week range will average to 48 weeks, but the additional weeks that agents are infectious while sub-clinical will push the number of infectious weeks to an average of greater than one year.

While these parameters are sampled on the first day of an agent's infection, they are not set in stone and can be impacted by treatment. For example, in default scenarios all sub-clinical or symptomatic agents will become non-infectious in two weeks and be cured in six months if they take treatment. Whether or not an agent takes treatment depends on whether the TB is diagnosed. Diagnosis in turn is impacted by interventions such as contact tracing and targeted universal testing.

Agents can be reinfected after having been cured, but the disease course of the second infection is not impacted by the first and the model does not include any form of TB immunity due to prior infection. The current version of the model also does not include TB relapse, drug-resistant forms of TB, or TB preventive therapy. Agents in our model also can't oscillate between sub-clinical and active TB.

While the traditional view is that people become non-infectious after two weeks on treatment (our default parameterisation), emerging evidence suggests a more complex picture by which people may remain infectious for much longer. (79) Accordingly, we included an alternative discrete probability distribution, indirectly derived from a review article published in 2021, (79) from which the model samples to determine when an agent will become non-infectious. The review estimated whether people on treatment would produce infectious sputum at various time-points, but explicitly did not take into account whether or not people stop coughing. This made direct use of the review findings unworkable for our purposes since we are interested in how infectious and agent is, which is determined both by whether sputum is viable and by whether someone is coughing. We thus, using the review as a rough guide, constructed an alternative distribution whereby agents become non-infectious anything from two weeks to 12 weeks after starting treatment. While this alternative distribution is thus in a sense arbitrary, we nevertheless consider it a useful tool with which to explore what impact longer infectious periods may have on the various epidemiological dynamics of interest. (figure 2B illustrates the default and alternative probability distributions for determining when agents will become non-infectious.)

2A

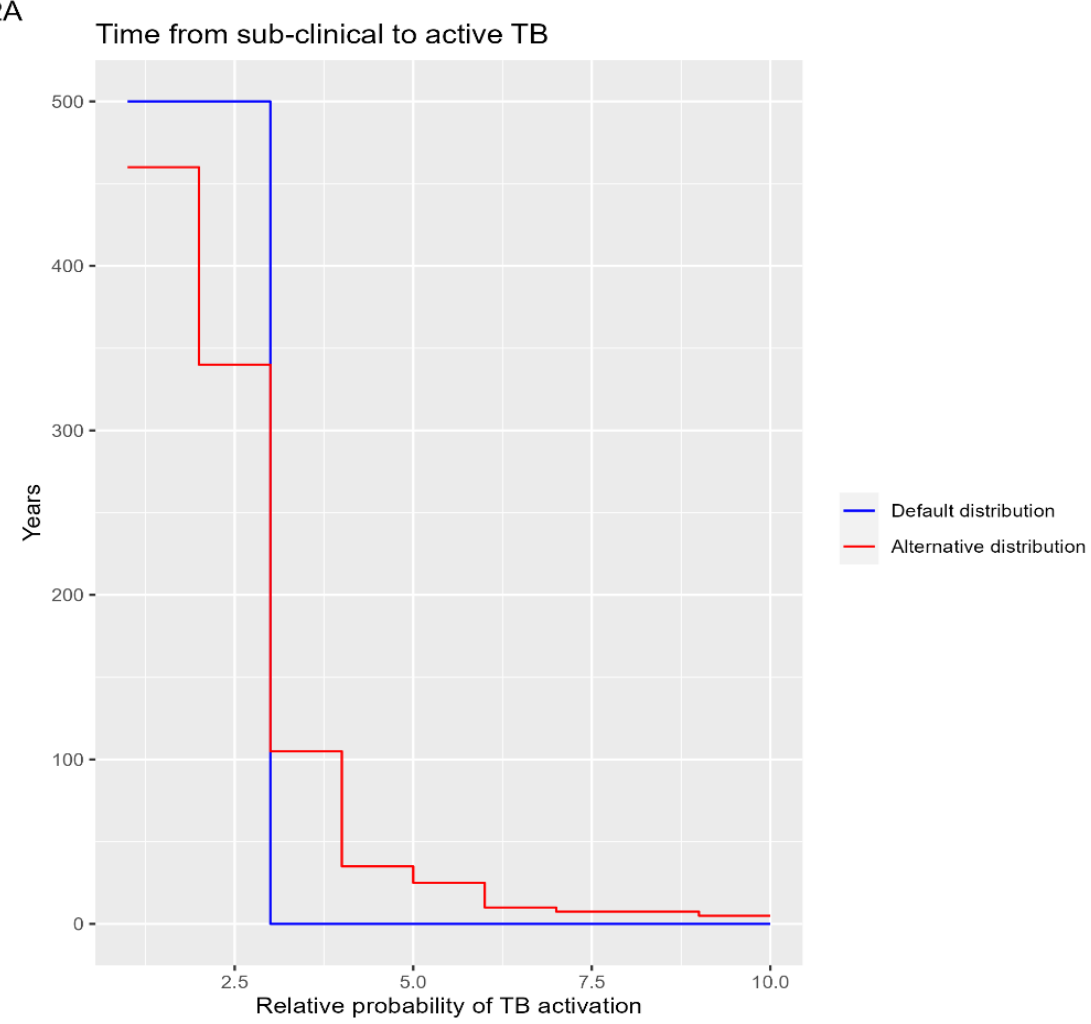
Time from sub-clinical to active TB
500
400
300
200
100
0
Years
2.5
5.0
7.5
10.0
Relative probability of TB activation
Default distribution
Alternative distribution

2B

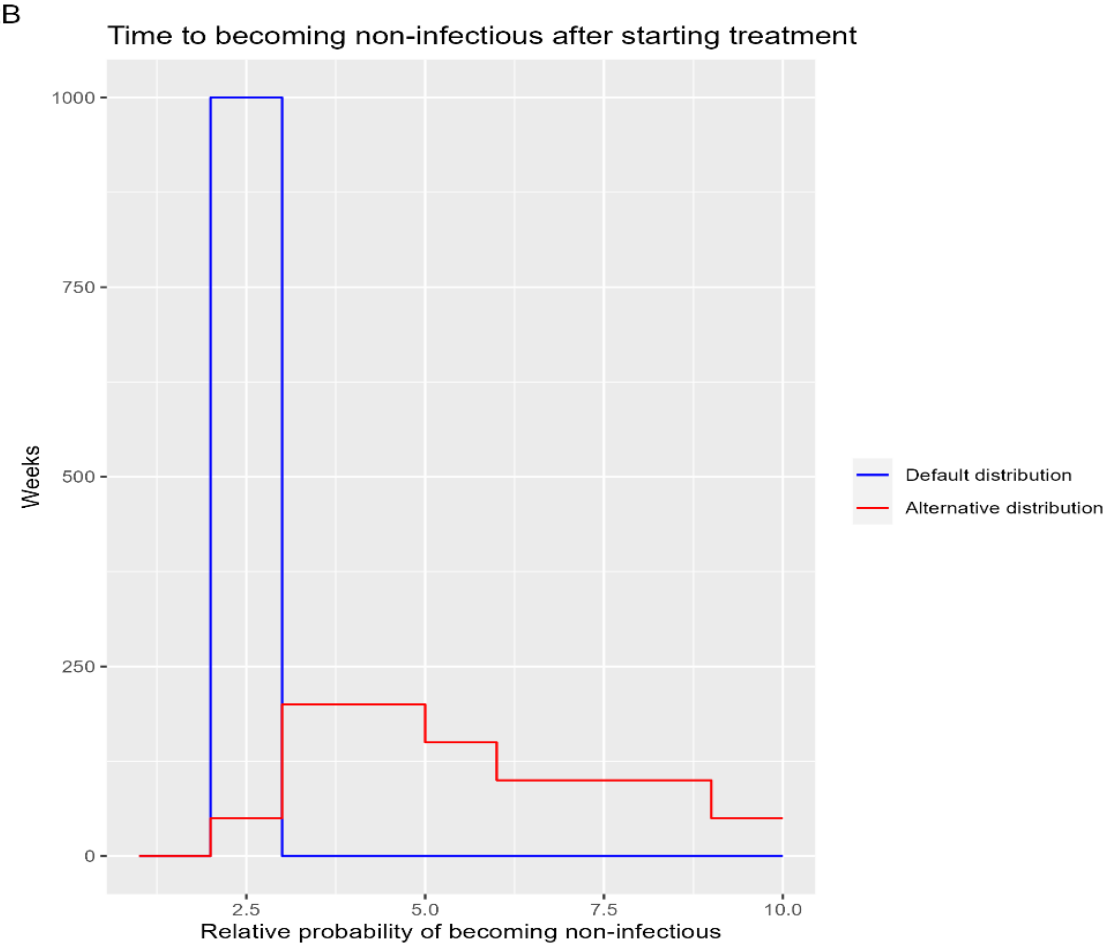

Figure 2: Alternate probability distributions for key parameters

Caption: These are probability distributions showing the relative likelihood of events occurring over time. In each graph the volume/area under the two lines are equal. Figure 2A shows time from becoming sub-clinical to becoming symptomatic and figure 2B shows time from starting TB treatment to becoming non-infectious.

Table 2: Disease progression parameter distributions in the TB model

Parameter	**Distribution/likelihood**
Likelihood of progression from latent infection to sub-clinical disease	10% (77)
Timing of progression from latent infection to sub-clinical disease	Four to 72 weeks uniformly distributed
Likelihood of progression from sub-clinical to symptomatic disease	100%
Timing of progression from sub-clinical to symptomatic disease	Four to 48 weeks uniformly distributed
Likelihood of progression from symptomatic	50% (78)

disease to death	
Timing of progression from symptomatic disease to death	24 to 72 weeks uniformly distributed
Likelihood of progression from symptomatic disease to cure (without treatment)	50%
Timing of progression from symptomatic disease to cure (without treatment)	24 to 48 weeks uniformly distributed

Calibration of transmission parameters

As mentioned in section d above, whether or not an agent becomes infected with latent TB, that is to say the transmission risk a specific agent is exposed to on a specific timestep in a specific setting, depends on two factors: the number of infectious agents associated with that space, and the inherent risk parameter associated with the type of space. The five inherent risk parameters, one for each type of space, are thus the key variables by which the transmission processes in the model can be calibrated. While the specific values of these parameters are not meaningful in themselves, together with the social interaction pattern determined by the model structure they control two key elements of the model: the overall force of infection and the proportion of infections that take place in each type of setting – in other words, how many infections take place and where they take place. (See chapter 3 for a discussion of the methodology and relevant formulas.)

Accordingly, the five inherent risk parameters were varied to simultaneously calibrate the model so as to (i) produce incidence rates within the range of incidence values that the WHO estimates for South Africa and (ii) to reproduce setting-specific transmission patterns matching those described by Andrews et al*.* in a study of TB transmission in a South African township. (71) The WHO estimates that TB incidence in South Africa in 2020 was 554/100,000 (range 388/100,000 to 749/100,000). (6) In the base scenario, our model is calibrated to produce a mean incidence over 10 years of 553/100,000. Using CO2 monitors together with social contact patterns, the Andrews et al*.* study estimated the proportion of infections by age group that take place in own households, other households, schools, workplaces and public transport in a Cape Town township. We combined these age-specific proportions to derive over-all proportions (for all ages) as follows: own household 15.6%, other households 10.9%, schools 10.1%, workplaces 41.4%, and public transport 21.9%. Our base model is calibrated on average to produce similar percentages. We can thus say with some confidence that our model produces roughly the same number of infections one would expect in the real world and that transmissions in the model generally occur in the types of settings where the available literature suggests they occur.

Since we are also interested in the impact of interventions on case detection rates in addition to TB incidence, we then calibrated the model so that in the base scenario case detection rates are similar to that estimated by the WHO for South Africa – thus establishing a case detection baseline. The WHO estimates treatment coverage of 58% (range 43% to 83%) for South Africa in 2020. (6) They

define treatment coverage as TB notifications divided by estimated incidence – which is interchangeable with the term case-detection rate, as used in our modelling. Accordingly, we varied the testing rate parameter (proportion of symptomatic agents who will test) in our model to produce a mean case detection rate of 58% in the base scenario. The testing rate parameter was then held stable at that level in the intervention scenarios.

Table 3 shows the above literature-based target values together with the mean values and variants produced by the 10,000 runs of the base scenario of the model used in the main analysis.

Table 3: Calibration results for the TB model

Calibration value	**Target value based on literature**	**Mean values produced by 10,000 runs of the model**
Cases of active TB per 100,000 per year	554 (range 388 to 749)	552 (461-634)
% infections in homes	15.6%	15.6%
% infections in schools	10.2%	10.1%
% infections in workplaces	41.4%	41.6%
% infections on public transport	21.9%	21.9%
% infections in other households/blocks	10.9%	10.9%
Case detection rate as %	58% (range 43% to 83%)	58%

Note: The range around the WHO point estimate for case detection is remarkably wide. This is a product of the underlying uncertainty regarding how many undetected TB cases there actually are in South Africa. In our model world we know exactly how many infections there are, and accordingly the IQR around the 58% case detection mean is very narrow at 57%-59% (what little variation there is, is driven purely by model stochasticity).

Finally, we opted not to calibrate our model to findings from South Africa's First National TB Prevalence Survey and the previously mentioned cluster-randomised trial of a TUT intervention. This is because the TB prevalence survey was cross-sectional (provided only a snapshot of a specific point in time) and we consider WHO estimates of prevalence in South Africa to be more reliable, in addition to being more recent. The TUT study was of very short duration compared to the period that we are modelling. These, and other differences between these studies and our modelling framework makes meaningful calibration difficult, if not impossible. In addition, calibrating our model to match the findings of these two studies would risk a type of over-fitting since we would be manipulating the very outcomes that the model is meant to shed light on. That said, we do provide a comparison of how our model estimates of the impact of TUT on case detection compares with that reported in the cluster-randomised trial.

The default timestep in the TB version of the model is one week and in the default scenarios the model is run for 520 weeks, or 10 years. Each scenario under investigation is run 10,000 times to account for the model's stochasticity. The key outputs are the mean cumulative number of TB cases, the mean cumulative TB mortality over 10 years, and the mean case detection rate over 10 years.

Results

Results of the main analyses

The mean incidence of active TB in the base scenario was 552/100,000 (range 461-634). As expected, mean incidence of active TB in all the intervention scenarios was lower – 536/100,000 (range 432-619) in the contact tracing scenario, 526/100,000 (range 433-605) in the mass X-ray screening scenario, 497/100,000 (range 394-581) in the TUT scenario, and 488/100,000 (range 387-582) in the ultra scenario. On most measures the ultra scenario did best, followed in order by the TUT, mass X-ray screening, contact tracing, and the base scenario. (See table 4 for details.)

The mean case detection rate in the base scenario was 58% (range 55.2%-61%). As expected, mean case detection rates were higher in all intervention scenarios – 59.4% (range 56.5%-62.5%) in the contact tracing scenario, 62.9% (range 59.8%-65.8%) for mass X-ray, 71.1% (range 67.8%-74.3%) for TUT, and 74.2% (range 70.8%-77.6%) for ultra.

The mean period for which an agent is infectious (sub-clinical period plus active TB period) is 50.3 (range 48.3-52) weeks in the base scenario. This period is slightly lower in the intervention scenarios – 49.8 (range 47.8-51.6) weeks in the contact tracing scenario, 49 (range 47.1-50.7) weeks in the mass X-ray screening scenario, 46.6 (range 44.5-48.4) weeks in the TUT scenario, and 45.8 (range 43.9-47.5) weeks in the ultra scenario.

Table 4: Key outputs for the five scenarios

Metric	**Base**	**CTI**	**X-ray**	**TUT**	**Ultra**
Mean TB cases/100,000	552	536	526	497	488
Mean case detection rate	58.0%	59.4%	62.9%	71.1%	74.2%
Latent TB infections	8,767	8,536	8,391	7,975	7,823
Active TB cases	1,105	1,071	1,052	995	977
TB deaths	274	259	240	205	190
Mean weeks	50.3	49.8	49.0	46.6	45.8

infectious (sub-clinical +active)					
Mean weeks active	24.5	24.6	24.4	23.2	22.7

Caption: This table shows key numbers for each of the intervention scenarios. Uncertainty ranges for these indicators are available in the paragraphs preceding the table.

Plotting the mean incidence of active TB over the 10 years for the five scenarios produces epidemiological curves that underline the relatively small differences between scenarios. (See figure 3.) Apart from most of the key numbers being somewhat lower in intervention scenarios, epidemics in these scenarios also tend to peak a few weeks later – base = week 326 (range 249-419), contact tracing = week 330 (range 252-433), mass X-ray screening = week 333 (range 250-425), TUT = week 333 (range 246-442), and ultra = week 338 (range 242-445).

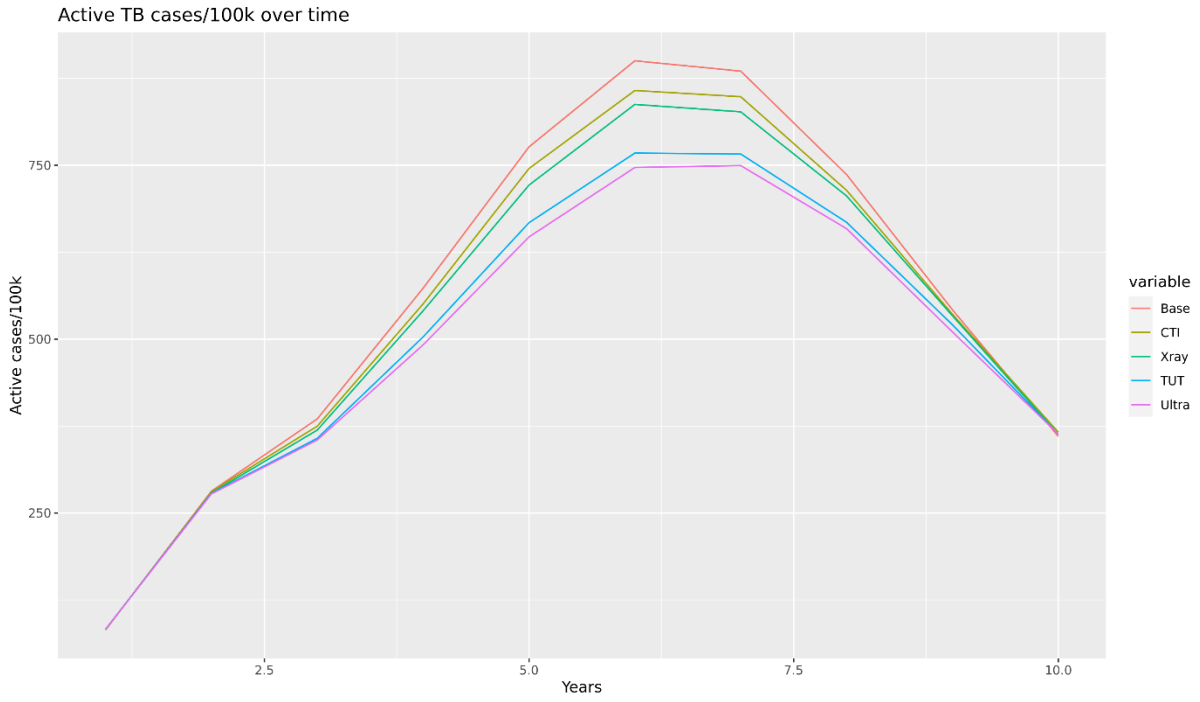

Figure 3: Epidemiological curves for the five scenarios

We also calculated the annual case detection rate for each of the 10 years in each of the five scenarios (detected new infections divided by actual new infections). As shown in figure 4, the case detection rate showed a slight, but steady increase over the 10 years.

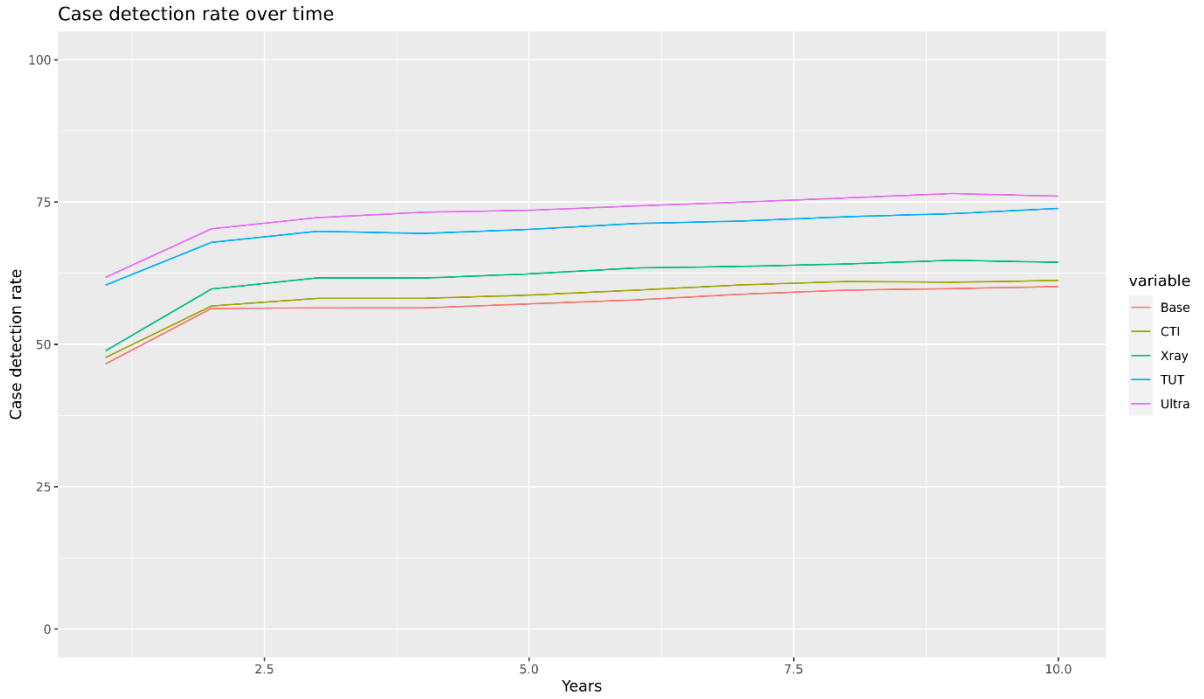

Figure 4: Case detection over time

As shown in table 5 below, the impact of all interventions was slightly greater on active TB cases than on latent TB infections, and much greater on TB deaths than on active TB cases.

Table 5: Impact of interventions compared to the base scenario

Metric	**CTI**	**X-ray**	**TUT**	**Ultra**
Latent TB infections	-2.3%	-4.0%	-9.0%	-10.8%
Active TB cases	-3.0%	-4.8%	-9.9%	-11.6%
TB deaths	-5.7%	-12.4%	-25.3%	-30.8%
TB cases/100,000	-3.0%	-4.8%	-9.9%	-11.6%

Caption: This table shows the reductions in mean values for key metrics compared to the base scenario.

Impact of varying key parameters and distributions

Given some of the uncertainties described in chapter 3, the model was also run with some alternative parameters and drawing from alternative distributions. While this does not amount to a full sensitivity analysis, these tests do give some insight into the relevant dynamics and were chosen specifically to anticipate potential blind spots in our model or areas where we thought the risk of uncertainty about key distributions skewing model outputs was greatest.

The trends identified above for the base configuration generally held for the alternative configurations, although the size and timespan of epidemics differed somewhat.

Firstly, instead of all agents becoming non-infectious after two weeks on treatment, as in the default configuration, we reran the model with the time to non-infectiousness being determined by sampling from a discrete probability distribution (as shown in figure 2A above). This resulted in a slightly larger epidemic, reflected in marginally higher TB incidence, latent TB infections, active TB cases, and TB deaths. (See table 6.)

Table 6: Impact of alternative probability distributions on key metrics

Metric	**Normal configuration (Base scenario)**	**Alternative non-infectiousness configuration (Base scenario)**	**Alternative TB activation configuration (Base scenario)**
Incidence (TB cases/100,000)	552 (461-634)	574 (483-654)	395 (279-498)
Latent TB infections	8,767 (7,939-9,469)	9,074 (8,244-9,751)	7,297 (5,915-8,359)
Active TB cases	1,105 (923-1,268)	1,148 (966-1,307)	789 (559-995)
TB deaths	274 (214-335)	287 (224-348)	182 (118-247)
Mean case detection rate	58% (55.2%-61%)	58% (55%-60.9%)	57.6% (54.2%-61%)

Secondly, we reran the model using the more complex distribution for latent-to-sub-clinical progression rather than the default uniform distribution (as shown in figure 2A above). This configuration resulted in somewhat smaller epidemics over the 10-year period being modelled – although the difference between this and other configurations would be reduced over longer time periods. As expected, this configuration resulted in later epidemic peaks – a mean of week 419 compared to a mean of week 326 with the base configuration. Interventions also tended to be slightly more effective in this scenario.

Table 7: Incidence reductions in the five scenarios using the alternate distributions

Scenario	**Normal configuration**	**Alternative non-infectiousness configuration**	**Alternative TB activation configuration**
CTI	-3%	-4.1%	-3.8%

X-ray	-4.8%	-6.4%	-6%
TUT	-9.9%	-11%	-13.1%
Ultra	-11.6%	-12.5%	-14.9%

Caption: This table shows the reduction in mean TB incidence for each intervention scenario compared to the base scenario in each of the three configurations.

Comparison with published data

Finally, we compared our model outputs to relevant data from the published literature. Having been calibrated to produce incidence and case detection rates that correspond to those estimated for South Africa, our model also produces death rates comparable to that in the literature – while we used risk of death parameters based on the published literature, we did not calibrate the model to reproduce the numbers of deaths seen in South Africa. The WHO estimates that there were 61,000 TB deaths in South Africa in 2020. With a population of roughly 60 million this amounts to just under 1 TB death/1,000 people. In our population of 20,000 agents (it is actually greater given births and natural deaths) one would thus expect roughly 200 deaths over 10 years. The base scenarios in our three configurations produced mean deaths of 274 (range 214-335), 287 (range 224-348), and 182 (range 118-247), respectively.

In a cluster randomised study of TUT in South Africa the intervention was found to result in a net increase in TB case detection of 17% per clinic per month compared to clinics where TUT was not implemented. In our primary analysis, TUT resulted in a mean increase in case detection over 10 years of 22.6%.

A key difference between our model outputs and the TUT study is that in our TUT scenario the intervention is implemented from the outset of the epidemic, whereas in the TUT study the intervention is implemented in an already mature epidemic. To allow for a more meaningful comparison, we accordingly ran our model with TUT implemented only from the beginning of year six. In these runs of the model, we found a mean improvement in case detection from year five to year six of 24.6% (or 14 percentage points). After this large jump from year five to six case detection remained relatively stable at the new higher level until year 10.

Since TUT is also impacted by HIV rates, we repeated this latter scenario with HIV prevalence set to 40% (roughly the maximum observed in communities in South Africa), rather than the 13% used in all our other scenarios. This resulted in a jump in case detection from year five to six of 62.4%, or 32 percentage points.

Discussion

In the ABM Spaces TB model, interventions aimed at the early detection of TB lead to small decreases in the number of latent TB infections, slightly greater decreases in the number of cases of active TB, and significantly larger decreases in the number of deaths caused by TB. These dynamics are not unexpected given that the type of interventions being modelled impact these key metrics in at least three ways – firstly by screening or testing agents who otherwise would not have screened or tested, secondly by screening or testing people who would have tested anyway but doing so at an earlier stage in their disease course, and thirdly by reducing the total number of infections over time. Whatever the complex and inter-related set of causes, the understanding that the same intervention can have a much greater impact on TB deaths than on TB case detection may be of use when evaluating case detection efforts in the real world.

The impact of the various interventions in our model was smaller than we expected, particularly on the number of latent TB infections and cases of active TB. We believe that this is largely due to the relationship between the timing of tests and the natural history and epidemiology of TB. Both mass X-ray screening and TUT are forms of screening and testing unrelated to whether or not an individual has any TB symptoms. These tests are thus in a sense randomly distributed in relation to the time at which agents become sub-clinical. If it is accepted that an agent can develop sub-clinical TB (i.e. become infectious) at any time, and is tested every 12 months, then on average one would expect agents to have been infectious for around six months before testing positive – in reality it is a bit more complicated since screening and testing coverage is below 100% and since some agents will test more quickly because they had become symptomatic. Either way, when looking for TB using annual screening or annual tests, our model suggests that a large number of agents would have already been infectious for some time prior to testing positive – and many of these agents would have transmitted TB in that time. We believe this dynamic explains the relatively low efficacy seen for annual interventions in our model.

Given the relevant time-scales, our findings could thus be interpreted as suggesting that more regular screening or testing than once a year of asymptomatic persons will be required in order to have a more substantial impact on TB incidence or mortality. While testing more frequently than once a year with X-rays or molecular TB tests are probably not realistic, it may well be that future urine or saliva-based screening tests could be done every month or every three months. In chapter 5 we build on these findings by exploring the impact of several levels of test frequency and test sensitivity on TB incidence and mortality.

Of all the interventions in our model, TUT was by far the most impactful. This is at least in part due to the fact that HIV prevalence in our model is high and that HIV increases the risk of progressing from latent TB infection to sub-clinical and active TB disease. In our model, as in the real world, TUT is by definition targeted at high-risk groups and this finding is thus not surprising.

The comparatively poor efficacy of mass X-ray screening in our analysis is because the screening was essentially untargeted and we've assumed far from universal coverage. As with TUT, the efficacy of mass X-ray screening will be highly context-specific – and it may, for example, perform much better

in comparison to TUT in communities with very high TB incidence and very low HIV prevalence. Indeed, as our additional analysis of the impact of TUT in communities with an HIV prevalence of 40% has shown, efficacy of these interventions can be highly dependent on the setting. One potential future use of ABM Spaces would be to more closely explore the impact of different levels of HIV prevalence on the impact of TUT – an analysis that could be used to inform cost-effectiveness calculations for TUT in different settings.

The impact of TUT and mass X-ray screening in our model was cumulative, that is to say that each continued to offer roughly the same additional benefit whether or not the other is also being implemented. This intuitively makes sense since, while there is some overlap in who will be reached by these tests, TUT will be more effective at specifically reaching people with HIV, while mass X-ray screening will be more effective at reaching asymptomatic HIV negative people. Either way, our findings support the concurrent use of these interventions.

The later epidemic peaks in the configuration with the alternative latent-to-sub-clinical progression was anticipated since agents in scenarios with this configuration will on average take longer to become infectious than in the base scenario – thus leading to slower or longer epidemics. The finding that interventions tended to be slightly more effective under this configuration, is likely because a slower, or stretched epidemic provides more opportunity for interventions to have an impact. The slightly larger epidemics in the alternative configuration for ‘time to non-infectiousness’ was also expected. The fact that the key trends regarding the effectiveness of the various interventions held in these alternative configurations suggests that these trends stand, irrespective of what the reality may be regarding the scientific uncertainties captured in the alternative configurations.

As with any modelling study, and particularly TB modelling, the work presented here has significant limitations. Most notably, there is uncertainty about various parameters relating to the natural history of TB. That uncertainty in the literature is unavoidably carried over into any modelling efforts. For the most part we have chosen parameters in line with current expert consensus, as reflected in the published literature – but it is worth stressing that significant uncertainty remains. In cases where the current or old consensus is being challenged, we have modelled both the old consensus and the new or alternative case – as for example with time from starting treatment to becoming non-infectious. As with the natural history of TB, there are also significant uncertainties relating to the epidemiology of TB that makes it hard to fit the model to real-world data. We have for example calibrated the model in part using WHO estimates for South Africa as a whole since sufficiently reliable community-level estimates could not be found.

Given these uncertainties, our model should not be seen as making precise predictions about TB in the real world. That said, while we caution against overinterpreting the specific numbers produced by this modelling exercise, we do think the structure and calibration of our model allows us to meaningfully explore the epidemiological dynamics relating to the various TB detection interventions. That is to say, the model roughly approximates real-world social networks, infections take place roughly where they take place according to previous studies, roughly as many agents become

infected as in the real world, and within this model world the various interventions are implemented in a way that is analogous to its real-world equivalent.

As with any modelling exercise, we have had to make choices about which processes to simplify and in which parts of the model to maintain complexity. This is inevitably a balancing act between maintaining nuance and capturing the key dynamics, on the one hand, and over-complicating the model to the extent of making it unworkable and impossible to calibrate, on the other. In the methodology section of this chapter and in chapter 3 we make the case that our choices are appropriate for our research questions.

Our model looks at one community with 20,000 agents at the outset. Even though there are births and natural deaths in our model, there is no movement of agents into or out of our population. This is different from real world dynamics whereby people often move in and out of communities, often bringing or taking infections with them. The epidemiological curves produced by our model (see figure 3) show a steeper ascent and decline than is seen in estimates of the TB epidemic in South Africa as a whole. This is as expected for three reasons: (i) in the real world different localised epidemics in different areas to some extent average each other out leading to flatter combined curves, (ii) epidemics in larger populations tend to take longer (the South African population is several orders of magnitude larger than our model population), and (iii) movement in and out of communities can extend epidemics. The fact that our model does not include TB relapse, means that incidence declines more rapidly than it might have done had relapse been included and counted as incident cases. Also notable in figure 3 is the fact that the difference between interventions is greatest when incidence is high (in the middle years) and converges as incidence declines toward year 10. This convergence suggests that there is a base level of transmission that none of the interventions modelled here can prevent – as hypothesised above, one factor contributing to this base level of transmission is that testing in our model is at best annual and thus most agents will have transmitted TB to at least some other agents prior to testing, irrespective of the intervention.

Finally, we did not attempt to do any cost effectiveness or cost estimates. Since all interventions modelled here involve doing additional tests, they will result in an increased direct cost of testing – at least in the short term. Over the longer term these interventions may well be cost-saving due to a reduction of TB incidence and hospitalisation, but that calculation is beyond the scope of this dissertation.

Chapter 5 – The impact of test frequency and sensitivity of a hypothetical TB test on TB incidence and mortality

Background

In chapter 4 we asked what impact several leading early case-finding strategies for TB can have on TB case detection, incidence, and mortality. One of our conclusions was that any testing strategy that relies on annual, or less frequent than annual, testing is likely to have only a limited impact on incidence. This is largely because by the time someone is diagnosed, they would on average have been infectious for around six months and would therefore already have transmitted TB to their close or regular contacts. Given these findings, an obvious follow-up question is what impact more frequent testing might have on incidence and mortality.

More frequent testing with the current gold-standard molecular tests is probably not feasible. This is because these tests require people to produce sputum – something many people find hard or uncomfortable. These tests also require the presence of a healthcare worker and the analysis of test samples is typically done at laboratories, although some point-of-care devices are also being developed. The WHO has published target product profiles to guide developers as to what to prioritise in the development of several types of TB diagnostics. (77-78) In particular, a report of a WHO meeting on such target product profiles set out several key targets that a TB diagnostic should meet – with high sensitivity and ease of use (not requiring sputum) being key factors. (17) Ease of use and low cost is likely to facilitate more frequent testing.

In this chapter, we contribute to the literature on target product profiles, or ideal future TB tests, by modelling the impact of a hypothetical test on TB incidence and mortality. By modelling 13 different levels of test frequency together with four levels of test sensitivity, we shed light on the relation between factors such as test frequency and incidence. In short, we ask how often we need to test if we are to see larger reductions in incidence.

Methodology

The TB version of ABM Spaces was adapted to model the use of a hypothetical TB test. Key elements of the model, including most of its structure and its parameterisation and calibration is identical to that presented in chapter 4. Population size (20,000) and length of simulation (10 years) were also the same. The only differences to the model presented in chapter 4 is that the TUT and mass X-ray routines were not used, a “regular testing” routine was added, and each scenario was run 1,000 instead of 10,000 times (explained below).

The model was run with testing frequencies of every four weeks, every eight weeks, every 12 weeks, and so on up to every 52 weeks (adding up to 13 levels of frequency). For each of the 13 levels of frequency, 1,000 simulations were run with test sensitivity set to each of 70%, 80%, 90%, and 100%. The analysis thus added up to a total of 52 scenarios and 52,000 simulations. Given this high number of scenarios with only very small differences between them, a total of 1,000 runs per scenario was considered sufficient for capturing the dynamics in question.

The levels of test sensitivity is in line with that discussed in a WHO meeting report on target product profiles which identified 98% as a minimum target for smear-positive TB and 68% as a minimum target for smear-negative disease. (17) ABM Spaces TB does not distinguish between smear-positive and smear-negative disease, but the range of values used in the model does cover a realistic set of average test sensitivity values. The WHO profile does not specify targets for test frequency, but it does place significant emphasis on ease of use and convenience, which we take to be preconditions for higher frequencies such as those that we model. Even so, our choice of four-weekly increments is somewhat arbitrary and was largely chosen as a mid-point between providing sufficiently granular information and managing the computational and analytical burden.

For each of the 52 scenarios we measured the mean number of TB cases, mean TB mortality, mean period for which agents are infectious, and the mean period for which agents have active TB disease (the infectious period is the period of active disease plus the period of sub-clinical disease). Our key question was whether the relationship between test frequency and these indicators are linear, and if not, whether there is an inflection point and where.

Results

In an adapted version of ABM Spaces, we found a clear correlation between more frequent testing and lower TB incidence and mortality. Higher test sensitivity was also correlated to lower incidence and mortality, but the impact was much less pronounced.

Four-weekly testing with test sensitivity of 100% resulted in a mean number of 361 (range 179-553) cases of active TB and a mean number of 35 (range 14-61) TB deaths over a ten year period in our population of 20,000 agents, which is 63.1% and 82.1% lower than the mean number of 978 (range 758-1,149) active cases and mean of 196 (range 142-244) TB deaths with 52-weekly testing and a sensitivity of 100%. With a test sensitivity of 70%, findings were similarly striking with mean active TB cases at 488 (range 261-695) and mean TB deaths at 51 (range 23-79) in the 4-weekly scenario compared to 1,012 (range 789-1,184) and 215 (range 154-265) in the 52-weekly scenario (reductions of 51.8% and 76.3%, respectively). See table 8 for more detail.

Table 8: Impact of test sensitivity on key indicators with four- and 52-weekly testing

Indicator	**70% sensitivity**	**80% sensitivity**	**90% sensitivity**	**100% sensitivity**
Mean active TB cases (four-weekly tests)	488 (261-695)	431 (214-640)	396 (216-603)	361 (179-553)
Mean active cases (52-weekly tests)	1,012 (789-1,184)	994 804-1,167	992 (794-1,162)	978 (758-1,149)
Mean TB deaths (four-weekly tests)	51 (23-79)	43 (17-72)	39 (17-65)	35 (14-61)
Mean TB deaths (52-weekly tests)	215 (154-265)	206 (153-257)	202 (148-250)	196 (142-244)
Mean infectious period in weeks (four-weekly tests)	30.5 (27.9-32.8)	29.2 (26.7-31.4)	28 (25.4-30.1)	27 (24.4-29.2)
Mean infectious period in weeks (52-weekly tests)	47.2 (45.1-48.9)	46.8 (44.6-48.5)	46.4 (44.4-48.1)	46 (43.9-47.8)
Mean active TB period in weeks (four-weekly tests)	17.1 (15.5-18.7)	16.5 (14.6-18.1)	16 (14.3-17.7)	15.6 (13.6-17.2)
Mean active TB period in weeks (52-weekly tests)	24.5 (22.8-26.2)	24.4 (22.6-26.2)	24.2 (22.6-25.9)	24.2 (22.4-25.8)

Caption: This table shows mean values for several key outputs for four-weekly and 52-weekly testing (the two extremes) for all four levels of test sensitivity. As elsewhere, the ranges are 95% uncertainty intervals.

The decrease in mean active TB cases and mean TB deaths associated with more frequent testing is not linear. As shown in figures 5 and 6, mean active TB cases and mean TB deaths declines relatively slowly as test frequency is reduced from 52-weekly to 16-weekly, whereafter it declines much more rapidly. Our modelling thus indicates an inflection point at around the 16-weekly (slightly under four-monthly) mark. (Note that the X-axis in these graphs do not represent time in years but represent different levels of test frequency from 4-weekly on the left to 52-weekly on the right.)

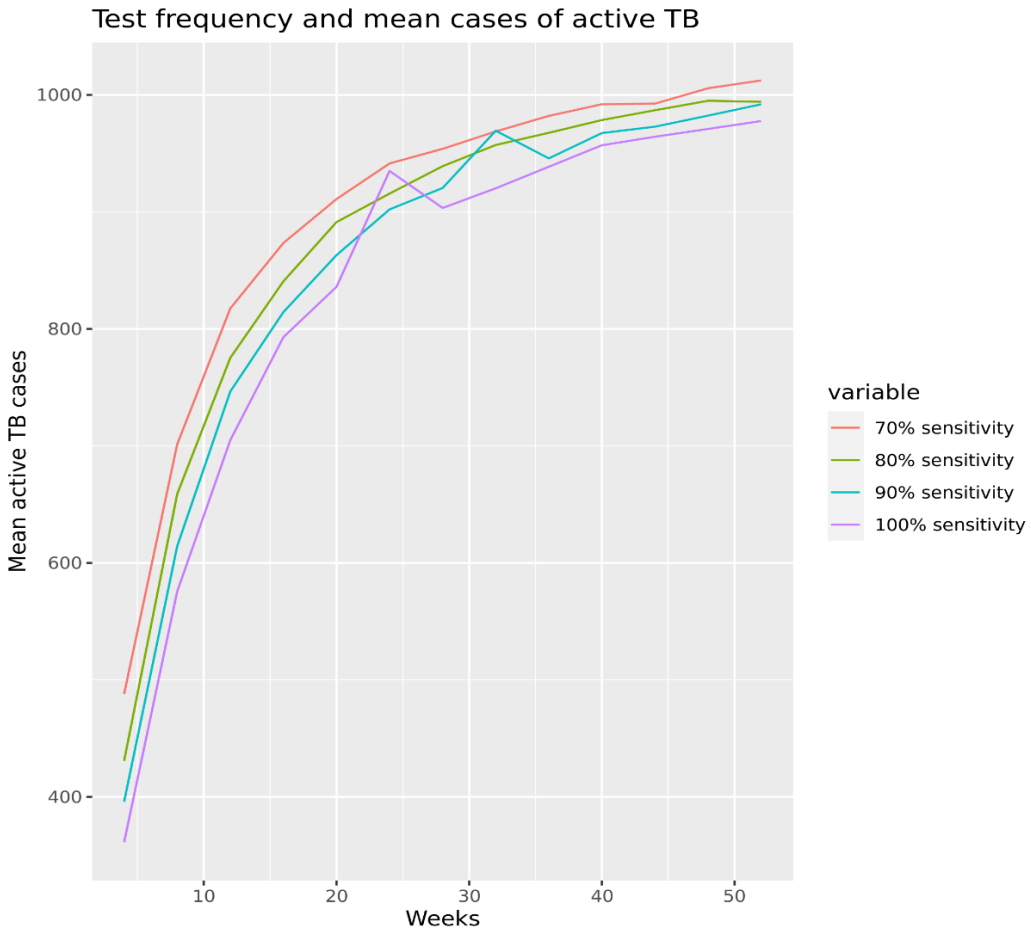

Figure 5: Test frequency and mean cases of active TB

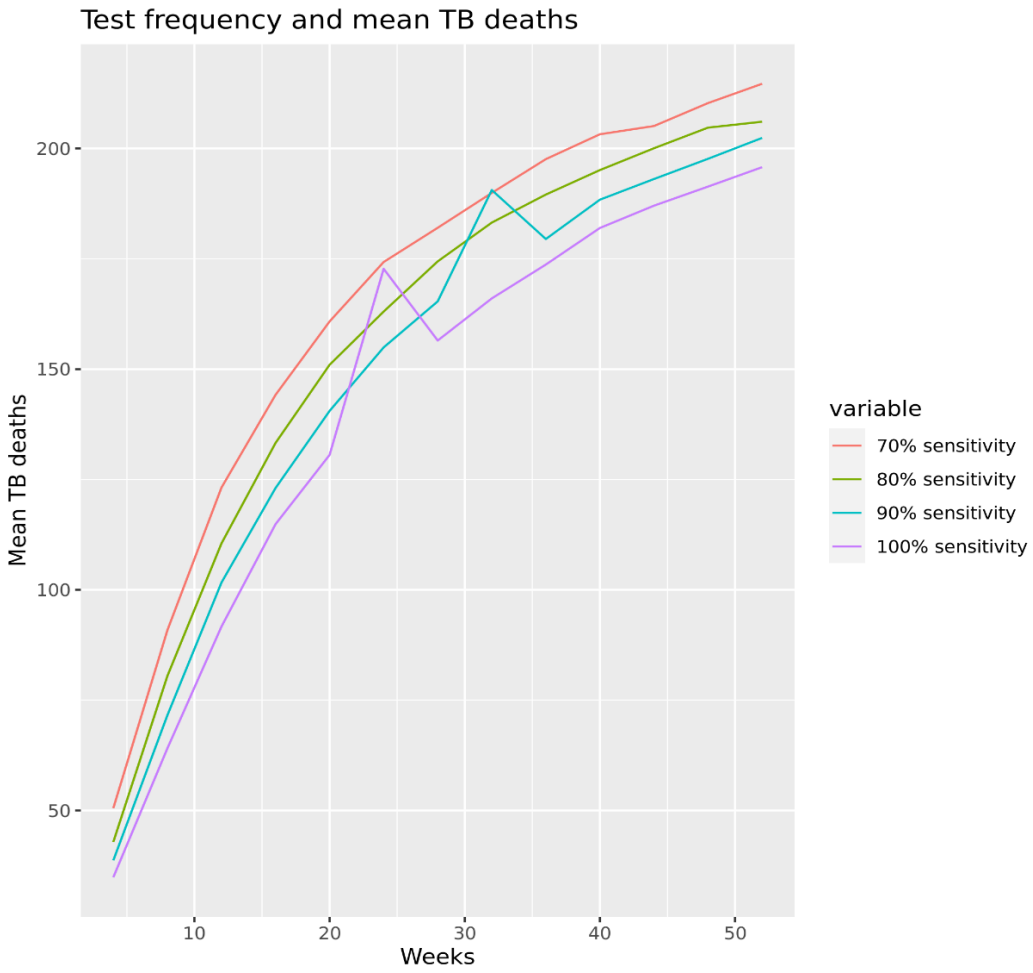

Figure 6: Test frequency and mean TB deaths

As shown in figures 7 and 8, there is a similarly non-linear pattern with an inflection point at around the 16-weekly mark for the mean period for which agents are infectious and the mean period for which agents have active TB disease. For example, from 52-weekly testing to 16-weekly testing in with test sensitivity of 90%, mean active TB cases falls by 178 from 992 to 814 – while from 16-weekly to four-weekly it falls by 418 (from 814 to 396).

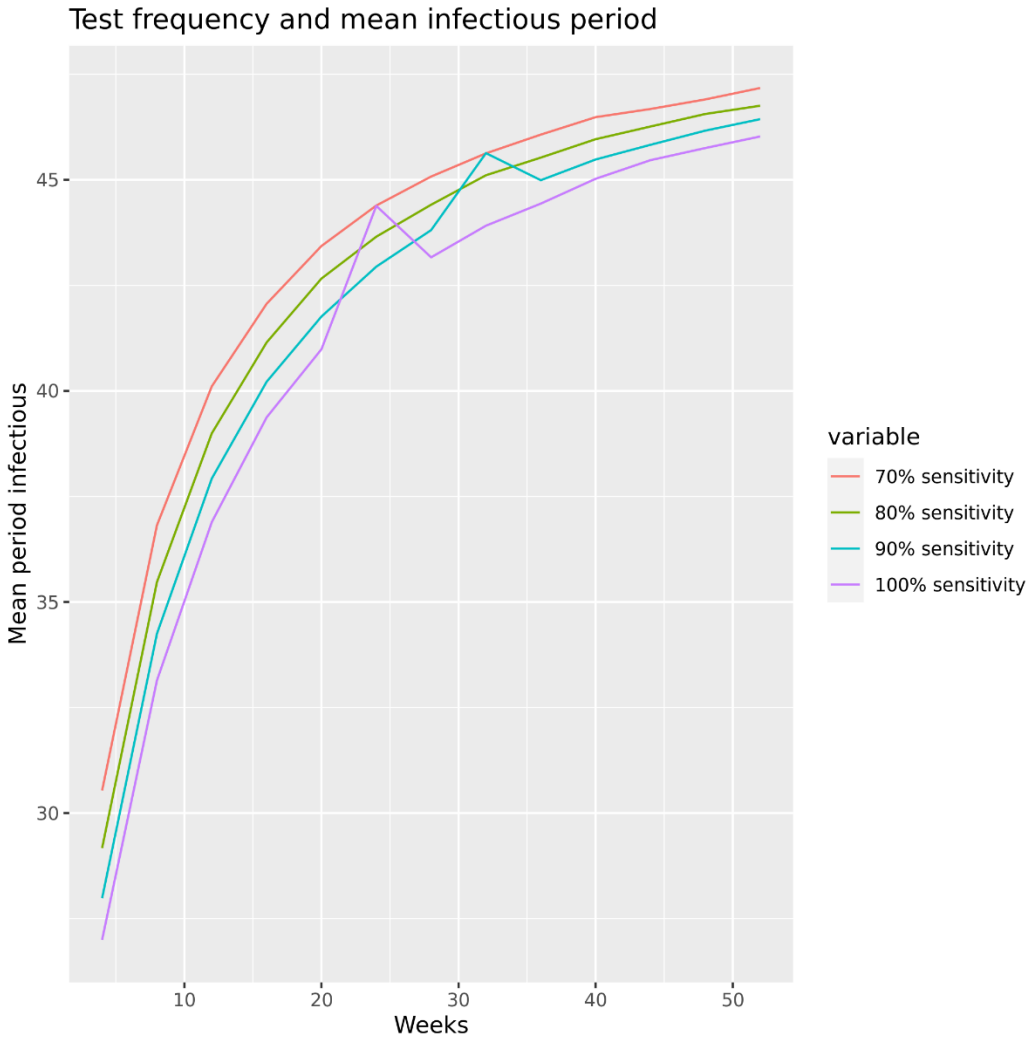

Figure 7: Test frequency and mean infectious period

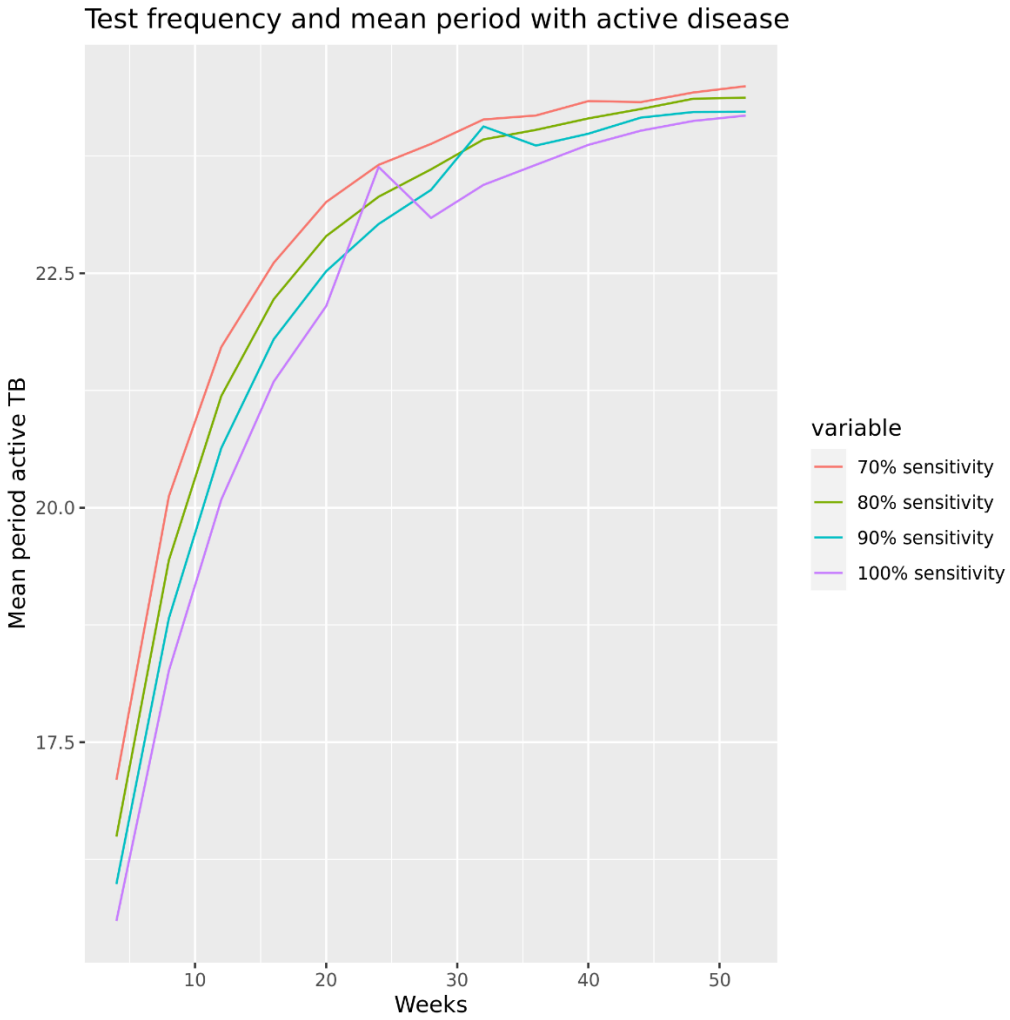

Figure 8: Test frequency and mean period with active disease

On the face of it these mean infectious and mean active disease periods (ranging from 27 to 47 weeks and 16 to 24 weeks, respectively) may seem longer than anticipated with high frequency testing. For example, four-weekly testing with a test sensitivity of 100% results in a mean infectious period of 27 weeks. There are two factors driving this effect. Firstly, agents do not immediately become non-infectious upon being diagnosed and typically only become non-infectious after two or more weeks on TB treatment (see chapter 4 for more detail on the default probability distribution for time from starting treatment to becoming non-infectious). Thus, even if all agents are diagnosed on the day that they become infectious, the mean infectious period cannot be less than two weeks. Secondly, even with highly frequent testing, not all agents are tested and some will slip through the cracks and remain infectious for longer periods, thus driving up the mean values. In addition, in scenarios where test sensitivity is lower than 100%, some agents who do test will not be diagnosed when tested and accordingly not be started on treatment.

Discussion

In an adapted version of ABM Spaces, we found a strong, non-linear correlation between increasing test frequency and reductions in the mean number of active TB cases and mean number of TB deaths in outbreaks. Higher test sensitivity was correlated with the same reductions, although the impact was much less pronounced. This finding that testing frequency has a greater impact on mean number of active TB cases and mean number of TB deaths than test sensitivity (at least in the test sensitivity range in our model) was not unexpected, although the scale of the difference was surprising. There is a high likelihood in the model that agents who receive a false negative test result will test positive in subsequent tests. In effect, a test sensitivity of 70% for a single test amounts to a much higher cumulative test sensitivity when that test is repeated every four weeks. It should be noted, though, that in our model agents are all exposed to the same risk of a false negative when test sensitivity is below 100% and the model thus does not account for the possibility that some people may in reality be more likely to get false negative test results than others, for example due to factors such as having a continuously low bacterial load.

Maybe most importantly, our modelling suggests that the relationship between testing frequency and TB incidence is not linear. Incidence declines relatively slowly as frequency is reduced from 12 months to around four months, and then starts declining much more rapidly. Additional research is required to establish where exactly this inflection point is. Either way, our model provides relatively compelling evidence in support of setting a target test frequency of three or fewer months. To get people to take a test that regularly, the test will have to be extremely easy to take – and would likely have to be a home test that makes use of either a urine or a saliva sample. Whether it is possible to produce affordable urine or saliva-based tests with sufficient sensitivity is not yet known.

The analysis presented in this chapter is unique in that it explores the impact of test frequencies as high as four-weekly. What literature we could find in this area typically looked at intervals of a year or longer between testing rounds. In line with our findings, a study published in PLOS ONE in 2011 found that more frequent testing results in more TB cases being averted.(82)

There are several limitations to the analysis presented in this chapter. Our hypothetical test is equally sensitive to sub-clinical and active TB, something that may not be the case with real-world tests. It is also possible that future real-world TB testing will involve multiple tests (for example to screen and confirm) rather than our ideal of a single definitive test. Our assumption of 100% specificity may also not be realistic. We assumed that on average around 80% of the population will regularly use the test. Such high usage may be possible in a high-prevalence community with a very cheap and really easy to use home urine or saliva test, but would be unlikely with more complex tests and in lower-prevalence settings. As in the previous chapter, we did not model the use of TB preventive therapy, nor did we include drug-resistant TB or the possibility of TB relapse. We also didn't model the monetary cost of more frequent testing. In real-world settings it is plausible that the concurrent use of TB preventive therapy would reduce the impact of highly frequent testing such as that modelled here.

Chapter 6 – The use of ABM Spaces to explore the impact of contact tracing and test turnaround times on COVID-19 infections and mortality

Note: ABM Spaces was initially conceptualised as a framework with which to model TB interventions. From the outset, however, the plan was to make the model easy to adapt for other infectious diseases where transmission is also mediated through environments like households and workplaces –which is the case when transmission is partly or mainly airborne. Accordingly, as COVID-19 was spreading across the world in early 2020, we pivoted the development of ABM Spaces first to address questions relating to COVID-19. This was both done because of the urgency of COVID-19 at the time, and because COVID-19 is in some respects simpler to model and thus presented an easier target for the first prototypes of ABM Spaces, before taking on the more complex challenge of modelling TB interventions. We include this chapter here to illustrate how the ABM Spaces model can be tailored to explore research questions in an infectious disease that differs from TB in important ways. An earlier version of the work presented in this chapter was published as a preprint on medRxiv in October 2020. (30) We have improved some of the original analysis and edited the paper to fit in this dissertation (for example by removing duplication of model description). We have not updated the model with new parameterisation and the analysis presented here is thus based on the original variant of SARS-CoV-2 and on what was known about the virus in mid-2020. We have added a note to the discussion section where we consider how these findings may change in the context of newer variants like Omicron.

Background

As discussed in chapter 2, as of mid-2020, many countries had implemented extended and far-reaching lockdowns in response to the COVID-19 pandemic, often at significant economic cost. A key question in the early months of the epidemic was what non-pharmaceutical interventions (NPIs) governments could implement in order to exit such lockdowns as safely as possible. (18) One NPI that received significant attention and that was endorsed by the World Health Organization was contact tracing and isolation (CTI). CTI involves tracing the recent contacts of a person who tested positive for SARS-CoV-2 and isolating these contacts, and in some cases also testing contacts.

One factor that can impact the efficacy of CTI is the amount of time it takes to get test results, referred to as the turnaround time (TaT). In the early months of the pandemic in South Africa the demand for tests far exceeded supply and testing capacity and TaT was often longer than 7 days.

Using the ABM Spaces model (parameterised and calibrated as described in Chapter 3) we explore the impact of CTI and variable TaTs on SARS-CoV-2 infections and mortality in the context of a high-density area such as a South African township, with interaction between individuals in households, schools, workplaces, public transport and within a neighbourhood block. The three main scenarios in

our modelling were (i) no CTI, (ii) CTI with testing of contacts, and (iii) CTI without testing of contacts (in both CTI scenarios contacts are isolated). To explore the impact of testing delays, each of these base scenarios were run with 10 different TaT values (two to 11) – adding up to a total of 30 scenarios. Each scenario was run 10,000 times, for a total of 300,000 simulations. For our primary analysis we focussed on the difference between TaT values of two and eight – although we report outcomes for all scenarios. We did not model any NPIs other than CTI.

Setup and calibration of the COVID-19 version of ABM Spaces

The COVID-19 version of the model was developed to simulate key dynamics of the transmission and disease progression of SARS-CoV-2 in a closed population of 10,000 agents loosely based on Masiphumelele, a high-density township in Cape Town, South Africa. Apart from its similarity to other high-density townships in South Africa, we chose Masiphumelele because of the significant amount of high-quality research that has been conducted on tuberculosis transmission in the township, some of which we have used to calibrate our model.

a. Agents in the COVID-19 model

Agents in the COVID-19 version of the model are given various attributes, including age, gender, and a variety of variables relating to COVID-19-relevant disease states (such as exposed, infectious, symptomatic, recovered, and dead). Though agents have an age attribute, agents in this version of the model do not age, given that the model is run for only 270 days. Age does however impact the model world by determining whether an agent is in school or eligible for work and by impacting an agent's mortality risk when ill with COVID-19.

b. Disease progression in the COVID-19 model

On the day an agent becomes infected with SARS-CoV-2, the model draws from a number of probability distributions and estimates of risk, based on the published Covid-19 literature, (79-80) to determine (i) whether and when the agent will become symptomatic, (ii) when the agent will become infectious, (iii) whether and when the agent will be tested, (iv) whether and when the agent will stay home, (v) whether and when the agent will require hospitalisation, and (vi) whether and when the agent will die. The model remains flexible regarding these values – for example, when hospitals are full an agent requiring hospitalisation will not be hospitalised and face a higher risk of death than would otherwise be the case. This latter feature means our model factors in the impact of limited ICU capacity on COVID-19 mortality.

The below flow chart (figure 9) shows the possible disease progression paths that an agent can take in the COVID-19 version of ABM Spaces.

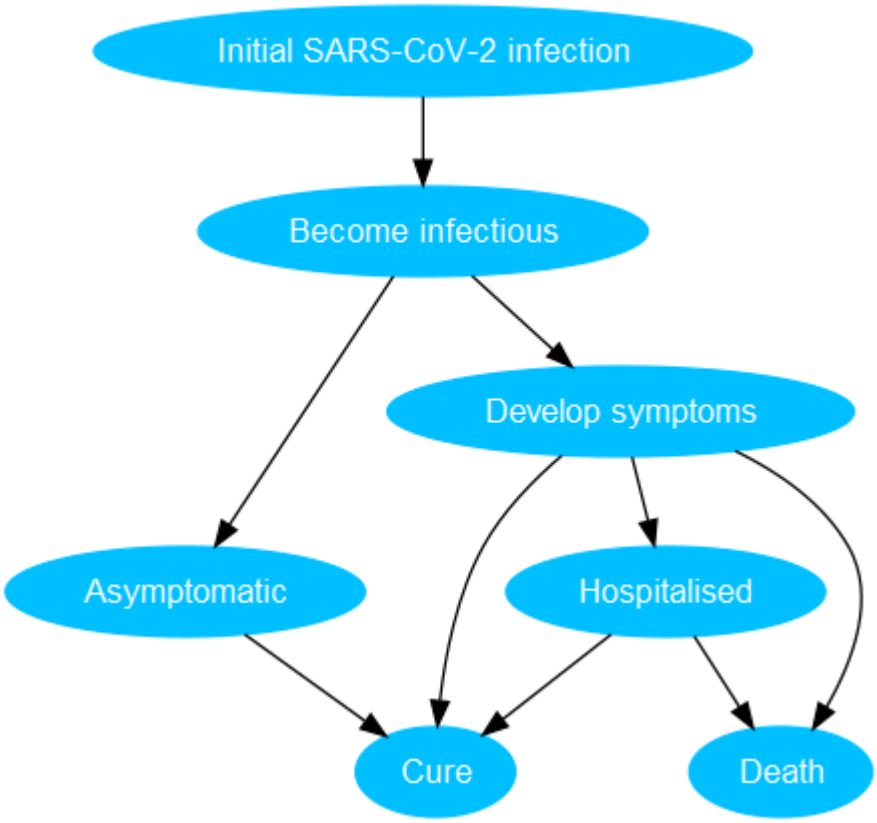

Figure 9: Flowchart of disease progression in the COVID-19 model

c. Spaces in the COVID-19 model

As in the TB version of the model, agents can be associated with five types of spaces (environments): households, workplaces, school classes, taxis, and city blocks. The choice of these settings is based on research conducted on tuberculosis transmission in Masiphumelele. (71) There was no such research available on COVID-19 transmission by type of setting available in mid-2020. Since TB transmission is also largely airborne, we considered these data to provide a close enough approximation of where COVID-19 transmissions might take place. Spaces in the COVID-19 version of the model function mostly precisely as they do in the TB version of the model. The only notable difference is that spaces are more stable since agents never age out of school classes or workplaces and since there are no births and natural deaths in the COVID-19 version of the model.

While the COVID-19 version of the model includes hospitalisation, hospitals in the model do not exist as a space where agents encounter each other and transmission takes place. This is because all agents in the model who are hospitalised are by definition already infected with SARS-CoV-2. A hospitalised agent is however withdrawn from all other spaces that it is associated with for the duration of its stay in hospital.

d. COVID-19 transmission

As in the TB version of ABM Spaces, transmission of SARS-CoV-2 occurs not directly from agent to agent, but from agent to setting to agent. On each timestep in the model, agents are exposed to the transmission risk of the various settings they are associated with. The setting and timestep-specific risk of a setting is a product of the number of infectious agents in that setting on that timestep and the inherent risk profile of the type of setting. As with the TB version of the model, the Reed-Frost formula is used to determine whether or not transmission takes place (the formula is discussed in chapter 3).

e. CTI and testing in the Covid-19 model

The model allows for false negative tests and takes into account published estimates of the variability of test sensitivity based on days since infection so that, for example, an agent with SARS-CoV-2 infection is much more likely to test positive on day four of infection as opposed to day one or day 14. (85) We gave symptomatic agents a likelihood of 0.8 of being tested in the days following first symptoms – although the model can easily be run with other values for this and other variables. The only way for asymptomatic agents to be tested in the model is as part of contact tracing. The model allows for different TaT values to be used. TaT values observed in South Africa ranged from two to greater than 11. (86)

The model can be run with or without CTI. In the default version of the model CTI involves testing 80% of an agent's household, school class or workplace contacts within two days of the agent testing positive. The test results for these tests become available a default of two days later (TaT equals two) whereupon the same process is repeated for all positive tests – with the modification that agents tested in the last 14 days will not be tested again. In CTI scenarios agents who test positive are isolated from all other agents and cannot further transmit the virus. Eighty percent of agents who are traced but who do not test positive, go into isolation (are temporarily removed from all spaces they are associated with), with the remaining 20% isolating at home (are removed from all their associated spaces apart from their household). The model can also be run so that a positive agent's contacts are isolated, but not tested (what we call the CTI without testing scenario).

Isolation refers to the movement of people with the infection being restricted, voluntarily or legally, while quarantine refers to the movement of people at risk of infection being restricted, voluntarily or legally. The work described here does not differentiate between these distinctions. We therefore use isolation to refer interchangeably to isolation or quarantine.

f. Calibration and parameter values in the COVID-19 model

Disease progress parameter distributions

In the first half of 2020, there was great uncertainty about the key epidemiological parameters relating to the natural history of SARS-CoV-2. Many studies were conducted to establish these parameters and efforts made to synthesise the available data. Here we refer largely to two sources. At the end of March 2020 Verity et al. published several key estimates in the *Lancet* medical journal, among others for the case fatality rate (1.38%) and the mean time from symptom onset to death (17.8 days). Their study also included estimates of case fatality rate by age, which were used to parameterise the case fatality rate for different age groups in our model. (83) Then in July 2020 the South African COVID-19 Modelling Consortium published a set of model projections including several parameters used in their modelling. (84) We consider these parameters to constitute expert consensus in South Africa at the time and accordingly drew on them as well. Table 9 below shows the key disease progression parameters used in the COVID-19 version of ABM Spaces.

Table 9: Disease progression parameter distributions in the COVID-19 model

Parameter	**Distribution/likelihood**
Timing of progression from exposure to becoming infectious	Two to six days with mean of four
Likelihood of progression to symptoms	25%
Timing of progression from becoming infectious to becoming symptomatic	One to three days with mean of two
Timing of progression from becoming symptomatic to becoming non-infectious	Four to six days with mean of five
Timing of progression from developing symptoms to testing	One to seven days with mean of four
Likelihood of requiring ICU admission	10%
Timing of progression from becoming symptomatic to ICU admission	Four to six days with mean of five
Case fatality rate	Overall 1.38%, but rates differ by age group as per Verity et al. (83)

Calibration of transmission parameters

In the COVID-19 version of the model, the parameters for inherent risk (*i*) associated with different types of spaces were calibrated using two processes: Firstly, setting-related risks were varied so as to reproduce published transmission patterns observed for TB in Masiphumelele. (10) Secondly, these

risks were calibrated so as to produce epidemic simulations with R_0 values that fall within the range of R_0 estimates found in the Covid-19 literature. (27)

We calculated R_0 values in the model using a numerical method that involves deriving R_0 from the equation $R = R_0 * S/N$ – where S is the susceptible population and N is the total population. (87) Since $R_t = 1$ at the peak of the epidemic and the fraction susceptible at the peak of the epidemic is known for each simulation, R_0 can accordingly be solved numerically.

As a result, the model is calibrated to produce a force of infection similar to that estimated in the literature with infections taking place in the settings where the available literature suggests they might.

Table 10: Calibration results for the COVID-19 model

Calibration value	**Target value based on literature**	**Mean values produced by 10,000 runs of the model in the base scenario**
R_0	1.5 to 3.5 (27)	1.68
% infections in homes	15.6%	15.8%
% infections in schools	10.2%	10.2%
% infections in workplaces	41.4%	41.4%
% infections on public transport	21.9%	21.7%
% infections in other households/blocks	10.9%	10.9%

g. Analysis in the COVID-19 model

The three base scenarios in the COVID-19 model were (i) no CTI, (ii) CTI with testing of contacts, and (iii) CTI without testing of contacts. To explore the impact of testing delays, each of these base scenarios were run with ten different TaT values (two to 11) – adding up to a total of 30 scenarios. Each scenario was run 10,000 times, for a total of 300,000 simulations. For our primary analysis we focussed on the difference between TaT values of two and eight – although we report outcomes for all scenarios.

All the above scenarios were run with populations constituted using published South African age structure data. To explore the impact of age structure, we ran an additional 10,000 simulations of the no CTI, TaT=2 scenario with a uniform age distribution rather than the South African age structure – by uniform here we mean that agents were no more likely to be one age than any other. Apart from the difference in age structure, all other parameters were kept the same for this additional analysis.

Each scenario was run for 270 timesteps (days) with a population size of 10,000 agents. Each scenario was run 10,000 times.

In the ABM Spaces model, the implementation of CTI results in significant benefits, particularly where contacts are tested, as shown in figures 10 and 11. CTI alone reduces the number of infections, but there is the greatest reduction with CTI with testing and a short TaT. Compared to no CTI, CTI with testing of contacts reduced the mean number of infections from 5,851 (range 5,692-6,013) to 4,439 (range 4,316-4,566) – a relative reduction of 24.1% – for a TaT of two days. The reduction was less pronounced for CTI only (without testing) – with a mean number of infections of 4,828 (range 4,694-4,963) and a relative reduction of 17.5%.

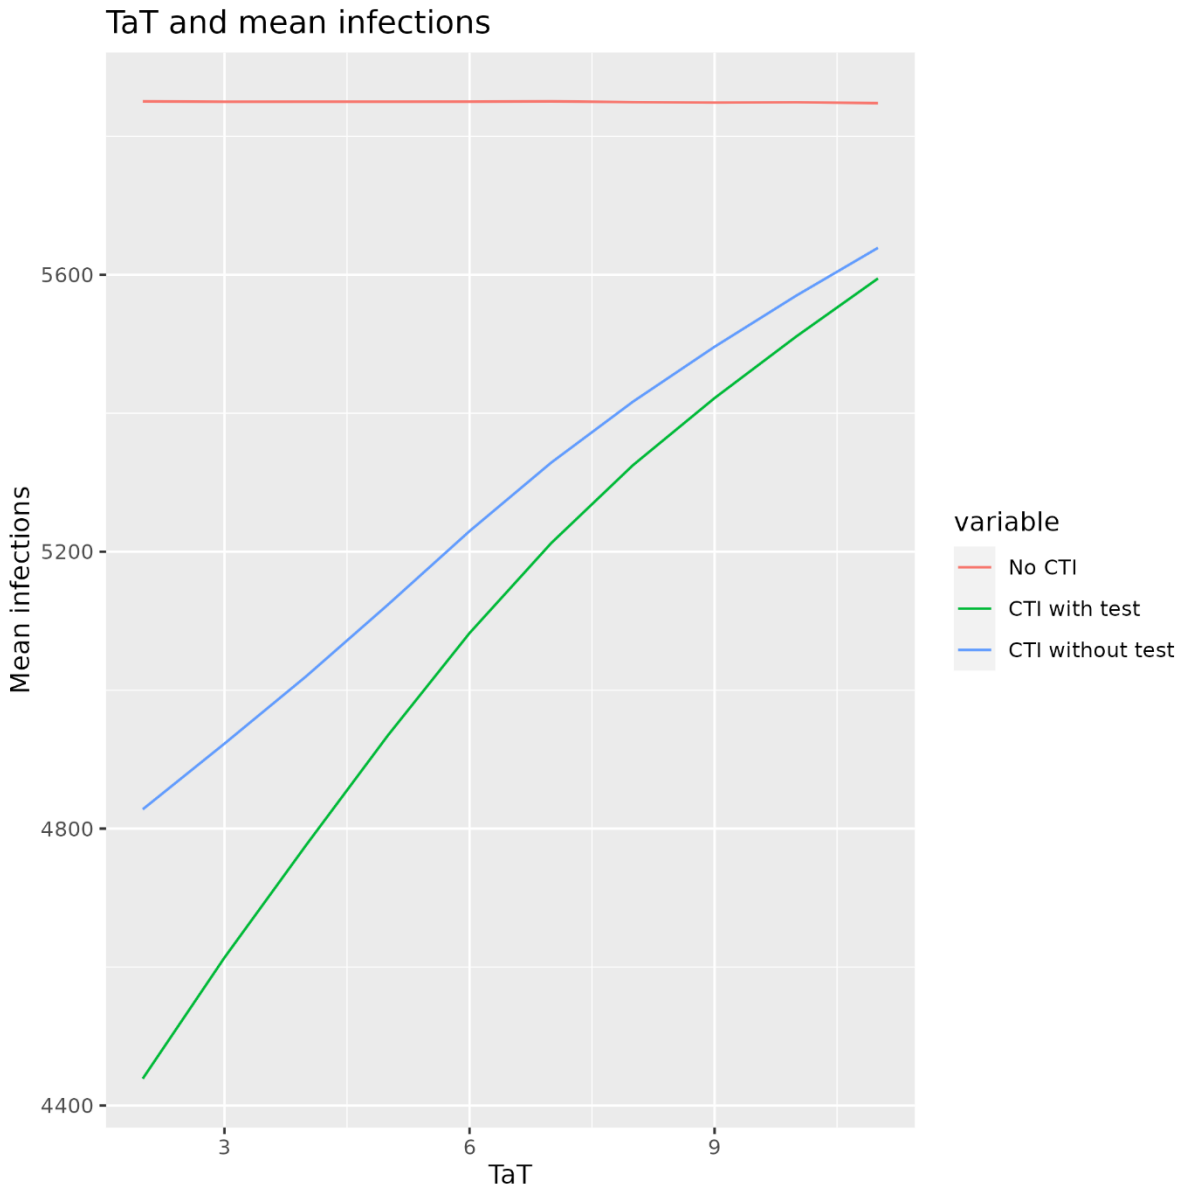

Figure 10: Test turnaround Time (TaT) and mean infections

CTI has a smaller relative impact on the number of deaths in the ABM Spaces model compared to infections. Again, the greatest reduction is observed with CTI with testing and a short TaT. For a TaT of 2 days, CTI alone showed a mean number of 33.3 (range 12-56) deaths – a relative reduction of 8.5% – whereas CTI with testing reduced the mean number of deaths from 36.4 (range 15-59) to 31.6 (range 11-54) – a relative reduction of 13.2%. Under this relatively ideal TaT, whether or not contacts were tested in the two CTI scenarios made only a moderate difference to deaths – with testing of contacts reducing mean number of deaths by 5.1%.

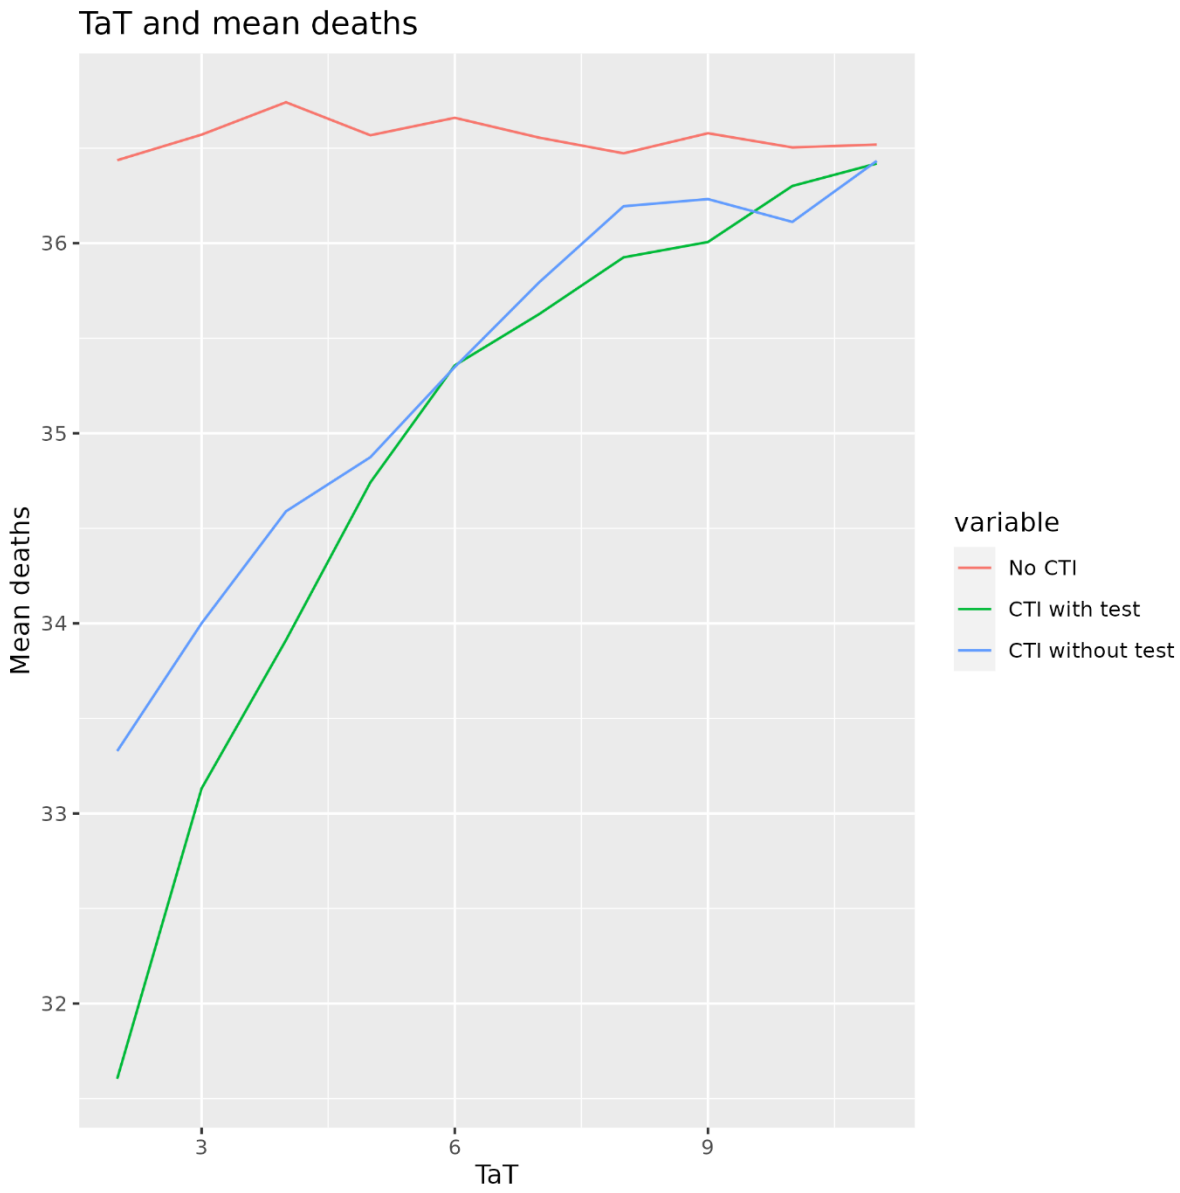

Figure 11: Test turnaround Time (TaT) and mean deaths

The testing of contacts, however, required dramatically more tests to be conducted. With TaT held at two days, a mean of 1,443 (range 1,369-1,519) tests were conducted in the no-CTI simulations, 11,870 (range 11,019-12,626) in the CTI with testing of contacts simulations, and 1,189 (range 1,126-1,252) in the CTI without testing of contacts simulations (the latter is lower than no-CTI since there

are fewer infections). In the two CTI scenarios and with TaT held at two days, one life was saved for every 6,283 additional tests conducted in the simulations where contacts were tested, and one infection averted for every 27.5 additional tests conducted in these scenarios.

While TaT had essentially no impact on mean infections in the no-CTI scenario, longer TaTs were associated with both more infections and more deaths in the two CTI scenarios. CTI with testing of contacts and a TaT of eight days resulted in a mean of 5,325 (range 5,184-5,468) infections and a mean of 35.9 (range 14-59) deaths – relative reductions of 9% and 1.6% compared to no CTI. CTI without testing of contacts and a TaT of eight days resulted in a mean of 5,417 (range 5,271-5,560) infections and a mean of 36.2 (range 15-59) deaths – relative reductions of 7.4% and 0.8% compared to no CTI. With a TaT of 8 days and no CTI, mean infections were 5,849 (range 5,689-6,011) and mean deaths were 36.5 (range 15-60).

The impact of both CTI and testing contacts on mean infections and mean deaths is greatest when TaT is two days – with the impact getting smaller as TaT increases. The difference between CTI with testing of contacts and CTI without testing of contacts also shrinks as TaT increases – a difference of 389 mean infections when TaT is two days drops to 92 when TaT is eight days.

In simulations with CTI and testing of contacts, the mean number of deaths was only moderately sensitive to TaT. In these scenarios the greatest difference was observed between low TaT values – mean deaths increased by 1.5 when TaT was three days compared to two, while an increase of only 2.8 deaths was seen with a TaT of eight days compared to a TaT of three. As with infections, the impact of TaT on the mean number of deaths was less pronounced in simulations without the testing of contacts but followed the same pattern of shrinking benefit.

Changing the age distribution to uniform (in the no-CTI, TaT=2 days scenario) resulted in mean infections of 4,829 (range 4,667-4,993) and mean deaths of 34.1 (range 15-55) – relative reductions of 17.5% and 6.3% compared to the same scenario with the South African age distribution. This finding is somewhat counter-intuitive since it has been hypothesised that South Africa's age distribution, with a youth bulge and few people over 60, would lead to fewer deaths, given that people over 60 are at much higher risk of dying of COVID-19 than younger people. In our model, however, having more young people (mixing in school classes) and working-aged people (mixing in workplaces) resulted in substantially more infections, and accordingly in more deaths.

Discussion

In an ABM that simulates key dynamics of Covid-19 transmission and disease progression, the WHO recommended strategy of CTI resulted in a small reduction in the mean number of deaths and a substantial reduction in the mean number of total infections. The benefit is greatest with TaTs of two days, but a marginal benefit remains even with TaTs of eight days or longer. In our model significant numbers of infections and deaths can be averted by testing the contacts of persons who test positive – a benefit that is greatest with shorter TaTs.

These results suggest that well-implemented CTI can play a valuable role in reducing the size of outbreaks in the real world and that TaTs should be kept as short as possible in order to maximise this benefit. Our findings thus broadly confirm those of Kretzschmar et al. and Hellewell et al. using a substantially different methodology. (25-26) Apart from the difference in methodology, our study differs from that of Kretzschmar et al. and Hellewell et al. in that we asked whether CTI and different TaT values lead to meaningful reductions in total infections, while they explored what level of intervention was needed to reduce the effective reproductive number to below one. It is possible that certain levels of intervention would substantially reduce total infections, but not so much as to bring the effective reproduction number below one.

This modelling work also illustrates a finding that would arguably be less likely with other, non-agent-based, types of modelling. Our model produced the counter-intuitive finding that the South African age structure would lead to more deaths than would be seen if age were uniformly distributed in the population. This finding is a product of the fact that most mixing in our model involves people of working or school-going age, something that produces more infections in the context of the South African age distribution than it does with uniformly distributed age. While our finding in this regard is far from conclusive, it does suggest a dynamic that was not anticipated.

Since several of our model parameters are highly uncertain, there is uncertainty in the model outputs in addition to the already significant stochasticity in the model. The model outputs should thus at most be considered as illustrative of underlying epidemiological dynamics rather than as an exact prediction of the impact of CTI and TaT in the real world.

The model has several limitations. Many of the Covid-19 epidemiological and disease progression parameters were highly uncertain in mid-2020 and remain relatively uncertain. The model does not account for co-morbidities such as diabetes, HIV and TB. Neither does it account for movement in and out of the community, a key element of Covid-19 transmission in the real world. Our model also does not include NPIs other than CTI, such as school closures. The setting-specific transmission-risk parameters in the model were deduced from published literature on TB transmission in a Cape Town community (to our knowledge no such South African data yet exists for Covid-19). Finally, we did not attempt to assess the cost-effectiveness of CTI in this study.

Note: The analysis presented here is based on the initial parent variant of SARS-CoV-2. More recent variants, such as Omicron, have shorter incubation periods and spread more quickly. We'd thus expect short test turnaround times to be even more important in the context of Omicron than they were in mid-2020. That said, the context has also changed in other ways. Most notably for this analysis, the development of rapid antigen tests means that test results can now be obtained within a few hours, rather than days – albeit at the cost of reduced sensitivity. The modelling of COVID-19 has also in some respects become more complex given that we now know that people can become infected multiple times and that much of the world's population have at least some form of immunity, be it natural or vaccine-induced – in mid-2020 it was still reasonable to model COVID-19 transmission without assuming the possibility of reinfection and the impact of vaccination. Either way, the

development of new SARS-CoV-2 tests, and arguably also the relative failure of contact tracing efforts in most countries, mean that the type of analysis presented here is less relevant in 2023 than it was in 2020. We nevertheless think it is worth including as an example of how a model like ABM Spaces can be used to explore urgent research questions.

Chapter 7 – Conclusions

After providing an overview of what ABMs are and reviewing their use in TB modelling, we used ABM Spaces to explore a set of specific questions relating to TB and COVID-19 detection and transmission. Based on this work, we can draw several broad conclusions – below we first note conclusions relating to the use of ABMs with substantial social structure and then conclusions relating to the specific TB and COVID-19 questions.

First, ABM Spaces illustrates that it is possible to develop ABMs with substantial social structure – in the form of multiple environments such as households, school classes, and workplaces, together with the complex social networks that link such spaces – in a way that is epidemiologically sound. Much of the TB ABMs we found in the literature contains only one or two such environments – ABM Spaces contains five and more could be added with relatively minor impact on model performance.

We described how model worlds where agents are associated with such spaces/environments in a relatively straight-forward manner can give rise to highly complex social networks. This type of modelling thus presents an alternative to models where complex social structure is hard coded or derived from sophisticated algorithms or geospatial data. Instead, we've shown that a relatively simple setup of letting agents be associated with specific households, workplaces or school classes, can give rise to rich and complex social structure in a way we think captures key aspects of real-world social networks without presupposing too much of the social structure. Such spaces/environments-based ABMs present an attractive trade-off between model complexity and computational burden that we'd encourage other modellers to consider when making decisions on type of model and model design.

Second, ABM Spaces provides evidence that complex ABMs with social structure can efficiently be coded in a compiled programming language such as C++, so that hundreds of complex simulations can be run in minutes on a personal computer. This shows both that access to supercomputers are not always required for this type of analysis and that compromises in terms of model complexity, as we've identified in the literature, are not always necessary. Epidemiological modellers who wish to make use of ABMs may thus benefit from greater collaboration with programmers skilled in compiled, high-performance languages such as C++. Our review of the literature suggests that such collaboration does not happen routinely and that modellers often use higher level modelling tools that are slower and less customisable. At the same time, our experience in adapting ABM Spaces to explore specific TB and COVID-19 questions, made it clear that significant epidemiological expertise and scientific rigour is required to produce robust models that are relevant to real-world epidemiological questions.

In short, our work suggests complex ABMs are best approached as inter-disciplinary projects requiring both scientific programming and epidemiological modelling expertise.

Third, while demographic data can be used to associate agents with different environments in TB ABMs (e.g. employment statistics can be used to determine how many agents should be associated

with workplaces), determining the different risk of disease transmission in these environments poses a challenge. Our novel solution to this problem was to calibrate ABM Spaces to published estimates of the places in which TB transmission takes place in a township outside Cape Town while also calibrating the force of infection in the model to match that estimated for South Africa. While we don't think that this calibration technique should be relied on for highly specific projections (given uncertainties in the data to which we are calibrating), we are confident that, by letting roughly the right number of transmissions take place in roughly the right types of places, the model sufficiently captures real-world transmission dynamics so as to allow for a meaningful exploration of epidemiological dynamics. Used cautiously, we believe such a calibration approach may be of value to other modellers.

Fourth, our analysis in chapter 4 found that the impact of several TB case-finding strategies on TB incidence, TB deaths, and TB case detection, was smaller than anticipated. Of the individual strategies, targeted universal testing was most impactful. The impact of combining strategies was additive, suggesting that in some cases it may make sense to implement multiple strategies simultaneously. The analysis supports the view that the choice of case-finding strategy or strategies is highly context dependent – with the case for targeted universal testing, for example, being much stronger in areas with high HIV rates. In a South African context, this would suggest that the case for TUT in a province with high HIV prevalence, such as KwaZulu-Natal, would be stronger than that in a province with lower HIV prevalence, such as the Western Cape.

Either way, our findings suggest that the South African government's decision to implement the TUT strategy was epidemiologically sound and is likely to lead to significant improvements in TB case detection, even if these effects may not be as large as we expected them to be. Our findings also support the wider use of X-ray screening in South Africa, as is already being done in several pilot projects, although it is not yet clear who should be offered X-ray screening, and how frequently, to achieve maximum impact.

While both these strategies offer individual and public health benefits, we did not model their cost-effectiveness and it is possible that in some contexts such case-finding strategies might not be cost-effective. Modelling the cost-effectiveness of such interventions in a way that accounts for the health benefits of earlier diagnosis and reduced transmission is a key area of future research.

Fifth, in interpreting why the TB case-finding interventions modelled in chapter 4 were less impactful in terms of reducing transmission than we expected, we reached the insight that any form of annual TB testing or screening of asymptomatic people would on average find people when they've already been infectious for at least six months (with the exception potentially of some blood marker tests that may have some predictive value). Since a substantial number of TB transmissions would already have occurred when someone has been infectious for six months, more regular testing may be required to make a significant impact on transmission.

Accordingly, in chapter 5 we explored the impact of a hypothetical TB test with several levels of test frequency and test sensitivity on TB incidence. We found that test frequency has a much greater

potential impact on incidence than test sensitivity (at least in the sensitivity range of 70% to 100% used in our analysis) – likely due to false negatives turning into true positives in subsequent cycles of testing. The reduction in incidence with more frequent testing in our model was not linear and showed a slow reduction from 12 months to around four months, and then a much more rapid decline down to four weeks. These findings suggest that a sufficiently sensitive test administered to roughly 80% of a community every three months, or even more frequently, if possible, could have a significant impact on the trajectory of the TB epidemic in that community.

Such regular testing with currently available tests is not feasible, but our work in this area does contribute to the growing literature on “ideal” TB tests as, for example, described in the WHO’s target product profiles. Monthly or three-monthly testing would require tests to be cheap and very easy and convenient to use. Very low-cost urine or saliva-based tests that can be self-administered at home would be ideal. Should such tests be developed, acceptability and cost-effectiveness studies will be an urgent priority.

While more frequent testing is one means to reduce the period for which someone is infectious, it is only part of the solution since, even if people start treatment immediately, they can remain infectious for several further weeks. Isolation of infectious persons is one means to reduce transmission in this period but may be particularly difficult to implement if people are asymptomatic or have only very mild symptoms, and isolation equates to time out of school or away from work. Tracing and testing contacts and providing TB treatment if positive and preventive therapy if negative (as per South African treatment guidelines published in 2023) (88) can also help to reduce further transmission. As with frequent testing, higher uptake of preventive therapy is also likely to depend on ease of use and convenience. Here also, as with TB diagnostics, new technologies such as a course of preventive therapy delivered as a single injection rather than a month or more of pills, could be transformative.

Either way, in the absence of a breakthrough new vaccine, the road to TB elimination is likely to be complex and involve a combination of many different new technologies. Our contribution here is to show that the potential gains from tests that can be administered or self-administered monthly or three-monthly could be substantial.

Sixth, the analysis presented in chapter 6 supports the intuitive view that for contact tracing to have a significant impact in controlling the spread of COVID-19, it is critically important that test turnaround times are as short as possible. The analysis was done based on data on the initial variant of SARS-CoV-2 – for later SARS-CoV-2 variants such as Omicron, where the incubation period is even shorter, test turnaround times would have to be even shorter to have the same impact.

More broadly, our modelling, as with much other COVID-19 modelling, confirms that contact tracing and isolation has the potential to be effective in reducing the spread of an infection. Translating this theoretical efficacy into real-world efficacy has however proven to be challenging – with maybe the exception of South Korea and a handful of other countries, contact tracing and isolation programmes in the COVID-19 pandemic have had very limited efficacy. This is partly due to the fact that COVID-19 presents a particularly tough challenge for contact tracing programmes since SARS-CoV-2 is a rapidly

spreading virus with substantial spread from asymptomatic people. There is however also an important reminder here that modelling complex social interventions such as contact tracing and isolation is difficult and there is always a risk of models oversimplifying complex and messy realities.

Seventh, one finding in chapter 5 illustrates how complex models such as ABMs with realistic social structure can produce emergent findings that may at first seem counter-intuitive, but that are nevertheless plausible. It was widely anticipated that South Africa's relatively young population would mean that comparatively fewer people per capita would die of COVID-19 in the country, given that COVID-19 is less deadly in young people than in old people. Yet, when we parameterised ABM Spaces with the South African age distribution (with a young population) we found a higher number of COVID-19 deaths than when we parameterised the model with a uniform age distribution (where agents were equally likely to be young or old). We were able to show that the reason for this was that young people in our model had more social contacts and there was thus more COVID-19 transmission, and accordingly deaths, in scenarios with the South African age distribution than in scenarios with a uniform age distribution. Our model does not provide any evidence that this is what happened in the real world, but it does illustrate a plausible epidemiological dynamic by which having a young population might actually result in more COVID-19 deaths rather than fewer. Our modelling here did not account for complicating factors such as the impact of lockdown measures such as school closures.

The capacity of ABMs to generate unanticipated outputs such as these may make them particularly valuable for the exploration of epidemiological dynamics. By letting agents interact in a realistic model world and seeing what emerges, the door is left open to more potential outcomes than when model dynamics are hard coded, as is essentially the case with models based on difference or ordinary differential equations. While the use of ABMs for intervention modelling is already quite common, there may well be a greater role for it in exploratory epidemiology than is currently the case.

Eighth, our work suggests that chance plays a large part in determining the speed and size of an outbreak in a specific community – while we reported mainly on the mean of 10,000 runs of the various scenarios, there was significant variation around those means. While this may simply be an artifact of the stochasticity in our model, a plausible case could be made that similar stochasticity drives outbreaks in real-world communities. In our model, for example, a high number of initial cases in a workplace that happens to be highly connected to households with many school-going children and adults associated with other workplaces could spark a rapid rise in infections – a dynamic that is also plausible in the real world.

One intriguing implication of this is that much of the variation seen in outbreaks between real-world communities may simply be due to chance. Such underestimation of the impact of random factors on the size of outbreaks in communities may also have implications for the interpretation of findings from cluster randomised trials. ABM Spaces could potentially be adapted to simulate cluster randomised trials in an attempt to better understand the impact of chance on such trials.

There are several other potential future uses of ABM Spaces. With relatively trivial modification, the model could be used to explore several other TB case-finding interventions not discussed in this dissertation. It could also be adapted for the exploration of early case-finding interventions for other infectious diseases where infection is mediated through spaces such as classrooms. Though a significant undertaking, a priority is to expand the TB version of ABM Spaces to include gender-related risk profiles, TB preventive therapy, TB relapse, drug-resistant forms of TB, and more variable levels of infectiousness.

Though we have made a case for the usefulness of ABMs, we are also very aware of their limits. None of the conclusions and findings in this chapter cite specific numbers. This is intentional. Complex ABMs like ABM Spaces contain many parameters and many internal complexities that together create a substantial risk of inaccurate model outputs with little relevance for the real world. Though this risk can to some extent be mitigated through careful calibration, there are simply too many variables and too many dynamics at play to allow much confidence in the precise numbers produced by such models, even if all possible precautions are taken.

That said, careful modelling such as that presented here, may nevertheless be sufficiently accurate to allow for the identification of general trends and indications of potential epidemiological dynamics that would be hard to identify using simpler models. We suspect that this, rather than precise prediction, is where the greatest value in complex ABMs is to be found. Indeed, with the examples in chapters 4, 5, and 6 we have illustrated that a specific type of complex ABM, one that contains substantial social structure in the form of spaces or environments, can shed light on important real-world epidemiological questions relating to the early diagnosis of TB and COVID-19. Though for most epidemiological questions compartmental modelling will remain the most appropriate choice, we have shown that for at least some problems, particularly those involving social interventions such as contact tracing, ABMs with substantial social structure offers a viable, and arguably superior, alternative.

Bibliography

1. Vynnycky E, White R. An Introduction to Infectious Disease Modelling. OUP Oxford; 2010. 401 p.

2. Cvjetanović B, Grab B, Dixon H. Epidemiological models of poliomyelitis and measles and their application in the planning of immunization programmes. Bull World Health Organ. 1982;60(3):405–22.

3. Johnson LF, May MT, Dorrington RE, Cornell M, Boulle A, Egger M, et al. Estimating the impact of antiretroviral treatment on adult mortality trends in South Africa: A mathematical modelling study. PLOS Medicine. 2017 Dec 12;14(12):e1002468.

4. van Leth F, van der Werf MJ, Borgdorff MW. Prevalence of tuberculous infection and incidence of tuberculosis: a re-assessment of the Styblo rule. Bull World Health Organ. 2008 Jan;86(1):20–6.

5. Pai M, Behr MA, Dowdy D, Dheda K, Divangahi M, Boehme CC, et al. Tuberculosis. Nature Reviews Disease Primers. 2016 Oct 27;2(1):1–23.

6. World Health Organization. Global Tuberculosis Report 2021 [Internet]. World Health Organization; 2021 Oct. Available from: https://www.who.int/publications/i/item/9789240037021

7. World Health Organization. Global Tuberculosis Report 2023. Geneva; 2023.

8. Calderwood CJ, Wilson JP, Fielding KL, Harris RC, Karat AS, Mansukhani R, et al. Dynamics of sputum conversion during effective tuberculosis treatment: A systematic review and meta-analysis. PLOS Medicine. 2021 Apr 26;18(4):e1003566.

9. Burke RM, Nliwasa M, Feasey HRA, Chaisson LH, Golub JE, Naufal F, et al. Community-based active case-finding interventions for tuberculosis: a systematic review. The Lancet Public Health. 2021 May 1;6(5):e283–99.

10. Drain PK, Bajema KL, Dowdy D, Dheda K, Naidoo K, Schumacher SG, et al. Incipient and Subclinical Tuberculosis: a Clinical Review of Early Stages and Progression of Infection. Clin Microbiol Rev [Internet]. 2018 Oct [cited 2021 Oct 11];31(4). Available from: https://journals.asm.org/doi/10.1128/CMR.00021-18

11. Achkar JM, Jenny-Avital ER. Incipient and Subclinical Tuberculosis: Defining Early Disease States in the Context of Host Immune Response. Journal of Infectious Diseases. 2011 Nov 15;204(suppl 4): S1179–86.

12. Scriba TJ, Fiore-Gartland A, Penn-Nicholson A, Mulenga H, Kimbung Mbandi S, Borate B, et al. Biomarker-guided tuberculosis preventive therapy (CORTIS): a randomised controlled trial. The Lancet Infectious Diseases. 2021 Mar;21(3):354–65.

13. Van der Walt M, Moyo S. Short Report: The First National TB Prevalence Survey, South Africa 2018 [Internet]. South Africa; 2021 [cited 2021 Oct 11]. Available from: https://www.knowledgehub.org.za/system/files/elibdownloads/2021-02/A4_SA_TPS%20Short%20Report_10June20_Final_highres.pdf

14. A CLUSTER RANDOMIZED TRIAL OF TARGETED UNIVERSAL TESTING FOR TB IN CLINICS [Internet]. CROI Conference. [cited 2021 Oct 11]. Available from: https://www.croiconference.org/abstract/a-cluster-randomized-trial-of-targeted-universal-testing-for-tb-in-clinics/

15. Marks GB, Nguyen NV, Nguyen PTB, Nguyen TA, Nguyen HB, Tran KH, et al. Community-wide Screening for Tuberculosis in a High-Prevalence Setting. New England Journal of Medicine. 2019 Oct 3;381(14):1347–57.

16. Ayles H, Muyoyeta M, Toit ED, Schaap A, Floyd S, Simwinga M, et al. Effect of household and community interventions on the burden of tuberculosis in southern Africa: the ZAMSTAR community-randomised trial. The Lancet. 2013 Oct 5;382(9899):1183–94.

17. High priority target product profiles for new tuberculosis diagnostics: report of a consensus meeting [Internet]. [cited 2023 Apr 12]. Available from: https://www.who.int/publications-detail-redirect/WHO-HTM-TB-2014.18

18. Han E, Tan MMJ, Turk E, Sridhar D, Leung GM, Shibuya K, et al. Lessons learnt from easing COVID-19 restrictions: an analysis of countries and regions in Asia Pacific and Europe. The Lancet. 2020 Sep;S0140673620320079.

19. Flaxman S, Mishra S, Gandy A, Unwin HJT, Mellan TA, Coupland H, et al. Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. Nature. 2020 Aug;584(7820):257–61.

20. World Health Organization. Contact Tracing in the Context of COVID-19: Interim Guidance 10 May 2020 [Internet]. Geneva; 2020 May [cited 2020 Jul 26]. Available from: https://apps.who.int/iris/rest/bitstreams/1277571/retrieve

21. Fisher D, Wilder-Smith A. The global community needs to swiftly ramp up the response to contain COVID-19. The Lancet. 2020 Apr;395(10230):1109–10.

22. Lee SW, Yuh WT, Yang JM, Cho YS, Yoo IK, Koh HY, et al. Nationwide Results of COVID-19 Contact Tracing in South Korea: Individual Participant Data From an Epidemiological Survey. JMIR Med Inform. 2020 Aug 25;8(8):e20992.

23. Walker P, Whittaker C, Watson O, Baguelin M, Ainslie K, Bhatia S, et al. Report 12: The global impact of COVID-19 and strategies for mitigation and suppression [Internet]. Imperial College London; 2020 Mar [cited 2021 Feb 17]. Available from: http://spiral.imperial.ac.uk/handle/10044/1/77735

24. Giordano G, Blanchini F, Bruno R, Colaneri P, Di Filippo A, Di Matteo A, et al. Modelling the COVID-19 epidemic and implementation of population-wide interventions in Italy. Nature Medicine. 2020 Jun;26(6):855–60.

25. Reiner RC, Barber RM, Collins JK, Zheng P, Adolph C, Albright J, et al. Modeling COVID-19 scenarios for the United States. Nature Medicine. 2021 Jan;27(1):94–105.

26. Kwok KO, Tang A, Wei VWI, Park WH, Yeoh EK, Riley S. Epidemic Models of Contact Tracing: Systematic Review of Transmission Studies of Severe Acute Respiratory Syndrome and Middle East Respiratory Syndrome. Computational and Structural Biotechnology Journal. 2019 Jan 1;17:186–94.

27. Hellewell J, Abbott S, Gimma A, Bosse NI, Jarvis CI, Russell TW, et al. Feasibility of controlling COVID-19 outbreaks by isolation of cases and contacts. The Lancet Global Health. 2020 Apr;8(4):e488–96.

28. Kretzschmar ME, Rozhnova G, Bootsma MCJ, van Boven M, van de Wijgert JHHM, Bonten MJM. Impact of delays on effectiveness of contact tracing strategies for COVID-19: a modelling study. The Lancet Public Health. 2020 Aug;5(8):e452–9.

29. Kucharski AJ, Klepac P, Conlan AJK, Kissler SM, Tang ML, Fry H, et al. Effectiveness of isolation, testing, contact tracing, and physical distancing on reducing transmission of SARS-CoV-2 in different settings: a mathematical modelling study. The Lancet Infectious Diseases. 2020 Oct;20(10):1151–60.

30. Low M, Geffen N. Contact tracing and isolation reduces covid-19 incidence in a structured agent-based model [Internet]. Epidemiology; 2020 Oct [cited 2021 Oct 13]. Available from: http://medrxiv.org/lookup/doi/10.1101/2020.10.06.20207761

31. Brauer F. Compartmental Models in Epidemiology. In: Brauer F, van den Driessche P, Wu J, editors. Mathematical Epidemiology [Internet]. Berlin, Heidelberg: Springer Berlin Heidelberg; 2008 [cited 2021 May 22]. p. 19–79. (Morel JM, Takens F, Teissier B, editors. Lecture Notes in Mathematics; vol. 1945). Available from: http://link.springer.com/10.1007/978-3-540-78911-6_2

32. Walker PGT, Whittaker C, Watson OJ, Baguelin M, Winskill P, Hamlet A, et al. The impact of COVID-19 and strategies for mitigation and suppression in low- and middle-income countries. Science. 2020 Jul 24;369(6502):413–22.

33. Shakeel SM, Kumar NS, Madalli PP, Srinivasaiah R, Swamy DR. COVID-19 prediction models: a systematic literature review. Osong Public Health Res Perspect. 2021 Aug;12(4):215–29.

34. Railsback, S. Grimm, V. Agent-Based and Individual-Based Modeling: A PRACTICAL INTRODUCTION. Princeton and Oxford: Princeton University Press; 2012.

35. Hontelez JAC, Lurie MN, Bärnighausen T, Bakker R, Baltussen R, Tanser F, et al. Elimination of HIV in South Africa through Expanded Access to Antiretroviral Therapy: A Model Comparison Study. PLOS Medicine. 2013 Oct 22;10(10):e1001534.

36. Williams B. Elimination of HIV in South Africa through expanded access to antiretroviral therapy: Cautions, caveats and the importance of parsimony. 2014 [cited 2022 Sep 25]; Available from: https://arxiv.org/abs/1403.7104

37. Wilensky, U. NetLogo [Internet]. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston, IL; 1999. Available from: http://ccl.northwestern.edu/netlogo/

38. The AnyLogic Company. AnyLogic. 2000.

39. Willem L, Verelst F, Bilcke J, Hens N, Beutels P. Lessons from a decade of individual-based models for infectious disease transmission: a systematic review (2006-2015). BMC Infect Dis. 2017 Sep 11;17:612.

40. Kasaie P, Dowdy DW, David Kelton W. An agent-based simulation of a Tuberculosis epidemic: Understanding the timing of transmission. In: 2013 Winter Simulations Conference (WSC). 2013. p. 2227–38.

41. Begun M, Newall AT, Marks GB, Wood JG. Contact Tracing of Tuberculosis: A Systematic Review of Transmission Modelling Studies. PLoS One. 2013 Sep 4;8(9):e72470.

42. Kasaie P, Andrews JR, Kelton WD, Dowdy DW. Timing of tuberculosis transmission and the impact of household contact tracing. An agent-based simulation model. American journal of respiratory and critical care medicine. 2014;189(7):845–52.

43. Tian Y, Osgood ND, Al-Azem A, Hoeppner VH. Evaluating the Effectiveness of Contact Tracing on Tuberculosis Outcomes in Saskatchewan Using Individual-Based Modeling. Health Educ Behav. 2013 Oct 1;40(1_suppl):98S-110S.

44. Tuite AR, Gallant V, Randell E, Bourgeois AC, Greer AL. Stochastic agent-based modeling of tuberculosis in Canadian Indigenous communities. BMC Public Health. 2017 Dec;17(1):73.

45. Espindola AL, Varughese M, Laskowski M, Shoukat A, Heffernan JM, Moghadas SM. Strategies for halting the rise of multidrug resistant TB epidemics: assessing the effect of early case detection and isolation. International Health. 2017 Mar;9(2):80–90.

46. Goodell AJ, Shete PB, Vreman R, McCabe D, Porco TC, Barry PM, et al. Outlook for tuberculosis elimination in California: An individual-based stochastic model. PLoS One. 2019 Apr 9;14(4):e0214532.

47. Gomez GB, Dowdy DW, Bastos ML, Zwerling A, Sweeney S, Foster N, et al. Cost and cost-effectiveness of tuberculosis treatment shortening: a model-based analysis. BMC Infect Dis. 2016 Dec 1;16:726.

48. Knight GM, Gomez GB, Dodd PJ, Dowdy D, Zwerling A, Wells WA, et al. The Impact and Cost-Effectiveness of a Four-Month Regimen for First-Line Treatment of Active Tuberculosis in South Africa. PloS one. 2015;10(12):e0145796–e0145796.

49. Reddy KP, Gupta-Wright A, Fielding KL, Costantini S, Zheng A, Corbett EL, et al. Cost-effectiveness of urine-based tuberculosis screening in hospitalised patients with HIV in Africa: a microsimulation modelling study. The Lancet Global Health. 2019 Feb;7(2):e200–8.

50. Reddy KP, Denkinger CM, Broger T, McCann NC, Gupta-Wright A, Kerkhoff AD, et al. Cost-effectiveness of a Novel Lipoarabinomannan Test for Tuberculosis in Patients With Human Immunodeficiency Virus. Clinical Infectious Diseases. 2021 Oct 5;73(7):e2077–85.

51. Shrestha S, Chihota V, White RG, Grant AD, Churchyard GJ, Dowdy DW. Impact of Targeted Tuberculosis Vaccination Among a Mining Population in South Africa: A Model-Based Study. American journal of epidemiology. 2017;186(12):1362–9.

52. McCreesh N, Karat AS, Baisley K, Diaconu K, Bozzani F, Govender I, et al. Modelling the effect of infection prevention and control measures on rate of Mycobacterium tuberculosis transmission to clinic attendees in primary health clinics in South Africa. BMJ Glob Health. 2021 Oct 25;6(10):e007124.

53. Salvatore PP, Proaño A, Kendall EA, Gilman RH, Dowdy DW. Linking Individual Natural History to Population Outcomes in Tuberculosis. J Infect Dis. 2017 Dec 27;217(1):112–21.

54. Zwick ED, Pepperell CS, Alagoz O. Representing tuberculosis transmission with complex contagion: an agent-based simulation modeling approach. Med Decis Making. 2021 Aug;41(6):641–52.

55. Badmus O, Camorlinga S, Simpson O. Poverty and the emergence of tuberculosis: An agent-based modelling approach. In: 2016 IEEE-EMBS International Conference on Biomedical and Health Informatics (BHI). 2016. p. 561–4.

56. Zhang W, Liu S, Osgood N, Zhu H, Qian Y, Jia P. Using simulation modelling and systems science to help contain COVID-19: A systematic review. Syst Res Behav Sci. 2023 Jan;40(1):207–34.

57. Lorig F, Johansson E, Davidsson P. Agent-Based Social Simulation of the Covid-19 Pandemic: A Systematic Review. JASSS. 2021;24(3):5.

58. Borges ME, Ferreira LS, Poloni S, Bagattini AM, Franco C, da Rosa MQM, et al. Modelling the impact of school reopening and contact tracing strategies on Covid-19 dynamics in different epidemiologic settings in Brazil. Global Epidemiology. 2022 Dec;4:100094.

59. Germann TC, Smith MZ, Dauelsberg L, Fairchild G, Turton TL, Gorris ME, et al. Using an Agent-Based Model to Assess K-12 School Reopenings Under Different COVID-19 Spread Scenarios – United States, School Year 2020/21 [Internet]. Epidemiology; 2020 Oct [cited 2023 Feb 15]. Available from: http://medrxiv.org/lookup/doi/10.1101/2020.10.09.20208876

60. Bicher M, Rippinger C, Urach C, Brunmeir D, Siebert U, Popper N. Evaluation of Contact-Tracing Policies Against the Spread of SARS-CoV-2 in Austria – An Agent-Based Simulation [Internet]. Epidemiology; 2020 May [cited 2021 May 22]. Available from: http://medrxiv.org/lookup/doi/10.1101/2020.05.12.20098970

61. Pham TM, Tahir H, van de Wijgert JHHM, Van der Roest BR, Ellerbroek P, Bonten MJM, et al. Interventions to control nosocomial transmission of SARS-CoV-2: a modelling study. BMC Med. 2021 Dec;19(1):211.

62. Huang Q, Mondal A, Jiang X, Horn MA, Fan F, Fu P, et al. SARS-CoV-2 transmission and control in a hospital setting: an individual-based modelling study [Internet]. Infectious Diseases (except HIV/AIDS); 2020 Aug [cited 2023 Feb 15]. Available from: http://medrxiv.org/lookup/doi/10.1101/2020.08.22.20179929

63. Ying F, O'Clery N. Modelling COVID-19 transmission in supermarkets using an agent-based model. Simuunza MC, editor. PLoS ONE. 2021 Apr 9;16(4):e0249821.

64. Nguyen LKN, Howick S, McLafferty D, Anderson GH, Pravinkumar SJ, Van Der Meer R, et al. Evaluating intervention strategies in controlling coronavirus disease 2019 (COVID-19) spread in care homes: An agent-based model – CORRIGENDUM. Infect Control Hosp Epidemiol. 2022 May;43(5):685–685.

65. Woodhouse MJ, Aspinall WP, Sparks RSJ, Brooks-Pollock E, Relton C. Alternative COVID-19 mitigation measures in school classrooms: analysis using an agent-based model of SARS-CoV-2 transmission. R Soc open sci. 2022 Aug;9(8):211985.

66. Salmenjoki H, Korhonen M, Puisto A, Vuorinen V, Alava MJ. Modelling aerosol-based exposure to SARS-CoV-2 by an agent based Monte Carlo method: Risk estimates in a shop and bar. Pausata FSR, editor. PLoS ONE. 2021 Nov 22;16(11):e0260237.

67. Lopes PH, Wellacott L, de Almeida L, Villavicencio LMM, Moreira AL de L, Andrade DS, et al. Measuring the impact of nonpharmaceutical interventions on the SARS-CoV-2 pandemic at a city level: An agent-based computational modelling study of the City of Natal. Eslava-Schmalbach JH, editor. PLOS Glob Public Health. 2022 Oct 14;2(10):e0000540.

68. Stapelberg NJC, Smoll NR, Randall M, Palipana D, Bui B, Macartney K, et al. A Discrete-Event, Simulated Social Agent-Based Network Transmission (DESSABNeT) model for communicable diseases: Method and validation using SARS-CoV-2 data in three large Australian cities. Ndeffo Mbah ML, editor. PLoS ONE. 2021 May 21;16(5):e0251737.

69. Ng V, Fazil A, Waddell LA, Bancej C, Turgeon P, Otten A, et al. Projected effects of nonpharmaceutical public health interventions to prevent resurgence of SARS-CoV-2 transmission in Canada. CMAJ. 2020 Sep 14;192(37):E1053–64.

70. Hoertel N, Blachier M, Limosin F, Sánchez-Rico M, Blanco C, Olfson M, et al. Optimizing SARS-CoV-2 vaccination strategies in France: Results from a stochastic agent-based model [Internet]. Health Policy; 2021 Jan [cited 2023 Feb 15]. Available from: http://medrxiv.org/lookup/doi/10.1101/2021.01.17.21249970

71. Andrews JR, Morrow C, Walensky RP, Wood R. Integrating Social Contact and Environmental Data in Evaluating Tuberculosis Transmission in a South African Township. The Journal of Infectious Diseases. 2014 Aug 15;210(4):597–603.

72. Hazelbag CM, Dushoff J, Dominic EM, Mthombothi ZE, Delva W. Calibration of individual-based models to epidemiological data: A systematic review. PLoS Comput Biol. 2020 May;16(5):e1007893.

73. Platt D. A comparison of economic agent-based model calibration methods. Journal of Economic Dynamics and Control. 2020 Apr;113:103859.

74. Statistics South Africa. Mid-year population estimates [Internet]. 2020 Jul. Available from: http://www.statssa.gov.za/publications/P0302/P03022020.pdf

75. Geffen N, Scholz SM. How various design decisions on matching individuals in relationships affect the outcomes of microsimulations of sexually transmitted infection epidemics. PLOS ONE. 2018 Aug 29;13(8):e0202516.

76. Behr MA, Edelstein PH, Ramakrishnan L. Revisiting the timetable of tuberculosis. BMJ. 2018 Aug 23;k2738.

77. Vynnycky E, Fine PEM. The natural history of tuberculosis: the implications of age-dependent risks of disease and the role of reinfection. Epidemiology & Infection. 1997 Oct;119(2):183–201.

78. Tiemersma EW, van der Werf MJ, Borgdorff MW, Williams BG, Nagelkerke NJD. Natural History of Tuberculosis: Duration and Fatality of Untreated Pulmonary Tuberculosis in HIV Negative Patients: A Systematic Review. Pai M, editor. PLoS ONE. 2011 Apr 4;6(4):e17601.

79. Calderwood CJ, Wilson JP, Fielding KL, Harris RC, Karat AS, Mansukhani R, et al. Dynamics of sputum conversion during effective tuberculosis treatment: A systematic review and meta-analysis. PLoS Med. 2021 Apr 26;18(4):e1003566.

80. Target product profile for next-generation drug-susceptibility testing at peripheral centres [Internet]. [cited 2023 Apr 12]. Available from: https://www.who.int/publications-detail-redirect/9789240032361

81. Target product profiles for tuberculosis preventive treatment [Internet]. [cited 2023 Apr 12]. Available from: https://www.who.int/publications-detail-redirect/9789240010130

82. Dodd PJ, White RG, Corbett EL. Periodic Active Case Finding for TB: When to Look? PLOS ONE. 2011 Dec 22;6(12):e29130.

83. Verity R, Okell LC, Dorigatti I, Winskill P, Whittaker C, Imai N, et al. Estimates of the severity of coronavirus disease 2019: a model-based analysis. The Lancet Infectious Diseases. 2020 Jun;20(6):669–77.

84. South African COVID-19 Modelling Consortium. National COVID Epi Model - Updated on 24 July 2020 [Internet]. South African COVID-19 Modelling Consortium; 2020 Jul [cited 2020 Jul 28]. Available from: https://sacovid19mc.github.io/

85. Kucirka LM, Lauer SA, Laeyendecker O, Boon D, Lessler J. Variation in False-Negative Rate of Reverse Transcriptase Polymerase Chain Reaction–Based SARS-CoV-2 Tests by Time Since Exposure. Annals of Internal Medicine. 2020 May 13;M20-1495.

86. NICD. COVID-19 Testing Summary: South Africa Week 23 [Internet]. NICD; 2020 [cited 2020 Jul 28]. Available from: https://www.nicd.ac.za/wp-content/uploads/2020/06/NICD-COVID-19-Testing-Summary_-Week-23-2020-005.pdf

87. An introduction to infectious disease modelling – EMILIA VYNNYCKY and RICHARD G WHITE [Internet]. [cited 2020 Jun 20]. Available from: http://anintroductiontoinfectiousdiseasemodelling.com/

88. Department of Health, South Africa. Guidelines for the treatment of TB infection [Internet]. Department of Health; 2023 Mar [cited 2023 Apr 12]. Available from: https://sahivsoc.org/Files/Health_Latent%20TB%20Infection_2023_web.pdf