An Econometric Method for Estimating Population Parameters from Non-Random Samples: An Application to Clinical Case Finding

Rulof Burger, Ph.D.
Associate Professor
Dept. of Economics
University of Stellenbosch
rulof@sun.ac.za

Zoë M. McLaren, Ph.D.*
Assistant Professor
School of Public Health
University of Michigan
zmclaren@umich.edu

DRAFT – Please do not cite without permission

This version: November 25, 2016

Abstract

Background
The problem of sample selection complicates the process of drawing inference about populations. Selective sampling arises in many real world situations when agents such as doctors and customs officials search for targets with high values of a characteristic. We propose a new method for estimating population characteristics from these types of selected samples.

Methods
We develop a model that captures key features of the agent's sampling decision. We use a method of moments with instrumental variables and maximum likelihood to estimate the population prevalence of the characteristic of interest and the agents' accuracy in identifying targets. We apply this method to tuberculosis (TB), which is the leading infectious disease cause of death worldwide. We use a national database of TB test data from South Africa to examine testing for multi-drug resistant TB (MDR-TB).

Results
Approximately one-quarter of MDR-TB cases were undiagnosed between 2004-2010. The official estimate of 2.5% is therefore too low and MDR-TB prevalence is as high as 3.5%. Signal-to-noise ratios are lower for new patients than for repeat patients.

Discussion
Our approach is widely applicable because of the availability of routinely collected data and abundance of potential instruments. Using routinely collected data to monitor population prevalence can guide evidence-based policy making.

1. Introduction

* Corresponding author. Both authors contributed equally to this work. Names are ordered alphabetically.

The problem of sample selection complicates the process of drawing inference about populations when the characteristic of interest is only observed for a non-random subsample of the population. This type of selective sampling frequently arises in many real world situations when agents are tasked with searching for targets with high values of a particular characteristic of interest. For example, doctors decide to test certain patients for diseases based on observed symptoms or risk factors; tax authorities select firms and individuals for audits based partly on the content of tax returns; customs officials investigate suspicious travelers and shipments; police pull over erratic drivers; universities and firms interview candidates with the most promising resumes. In most of these cases, data is passively collected during the process, which provides an opportunity to draw inference about these populations if the sample selection problem can be addressed. In this study, we propose a new method for estimating population characteristics from a selected sample of observations drawn using the abovementioned "high-value search" sampling mechanism, which is intended to identify population targets with high-values of a particular characteristic, not facilitate the estimation of specific population attributes.

When it is not possible to perfectly discern high-value from low-value targets or to select all targets in the population for testing, agents must triage available resources to test only those targets deemed most likely to have high values of the characteristic of interest. We develop a model that captures key features of the agent's decision making surrounding the sampling of targets from the population. We then use a method of simulated moments and maximum likelihood to estimate (a) the prevalence of the characteristic of interest in the population and (b) the accuracy with which agents are able to identify high-value targets. In addition, we implement an instrumental variables method that leverages the contrast between exogenous, discontinuous changes in the availability of resources for testing and the gradual change of the population characteristics. These types of sample selection adjustments are an attractive alternative to random samples or testing the full population. Performing random sample surveys of the population can provide unbiased estimates, but these types of surveys are infrequently conducted because they entail significant financial and time costs. Testing or selecting all observations in the population is generally neither feasible nor cost-effective.

We apply our method to the problem of estimating disease prevalence. Accurate information on disease prevalence is essential for health policy making so that limited resources can be targeted globally and nationally to improve patient outcomes and maximize overall population health. We focus on tuberculosis (TB), which kills more than 1.5 million people annually and recently surpassed HIV to become the leading infectious disease cause of death in the world (WHO 2015). Though multi-drug resistant TB (MDR-TB) patients comprised an estimated 5% of TB cases notified in 2014, they accounted for 13% of TB deaths and 20% of TB spending worldwide (WHO 2015). Early and accurate diagnosis of MDR-TB is therefore critical. However, only 12% of incident TB cases were tested for MDR-TB. The under-detection of MDR-TB drives the development of new forms of drug resistance such as extensively (XDR) and totally drug-resistant TB worldwide (Klopper et al. 2013).

We use 7 years of data from South Africa's National Health Laboratory Service (NHLS) database, which has wide national coverage and includes individual patient TB test results. Our identification strategy leverages a series of national and local policy changes that increasingly prioritized MDR-TB as an urgent health concern. Our methods enable us to obtain precise estimates of the prevalence of drug-resistant TB in the population. We find that between 17 and 30 percent of all MDR-TB cases were undiagnosed between 2004-2010 in South Africa, which worsens patient outcomes, increases transmission and leads to the development of additional drug resistance. These estimates are validated against results from national surveillance surveys that use random sampling methods.

The contribution of this paper is two-fold. First, we develop a new method for estimating population parameters from non-random samples, which is widely applicable to many settings. We demonstrate the ease with which this method can be applied to other contexts because it uses existing data, is low cost and can be quickly scaled up. The HIV literature has demonstrated the importance of using statistical methods to adjust ante-natal care and population prevalence estimates for representativeness (see Sakarovitch et al. 2007, Nyirenda et al. 2010, Bärnighausen et al. 2011, Hogan et al. 2012, Clark and Houle 2014, McGovern et al. 2015), however we develop more rigorous methods for routine data and are the first to apply these types of methods to TB. Second, by applying the model to the case of MDR-TB, we are the first to demonstrate that approximately one-quarter of MDR-TB cases in South Africa went undiagnosed between 2004-2010, which is a significant threat to TB control. Our method can be employed in low- and middle-income countries to cost-effectively develop guidance for health policy making and ultimately improve population health.

2. Model

2.1 Assumptions

Consider a population characterized by two variables x and y, which have the joint cumulative distribution function $F_{X,Y}(x, y)$ such that $E(y|x)$ is a monotonic function of x, scaled to be non-decreasing. This assumption seems particularly reasonable in cases where the observable x variable is partly caused by the unobservable value of y. For example, if a disease is one of several causes of experiencing a symptom, then the probability that an individual has the disease increases with the severity of the symptom. Suppose the analyst is interested in estimating the (unconditional) population mean of y denoted $\mu \equiv E(y)$, but the analyst only has access to a non-random (i.e. selected or purposive) sample. This sample was drawn by agents with the purpose of selecting target observations with the highest expected values of y (unobservable at the time of sampling) based on values of x, which are observable to the agent but not necessarily to the analyst. Sampling reveals the values of y for the sampled observations, but has an associated cost.

We demonstrate how, providing there is sufficient exogenous variation in the proportion sampled and under specific assumptions about stationarity regarding population characteristics and the agents' sampling algorithms, we can draw inferences about both the sampling mechanism and the distribution of y in the population.

Define θ as the proportion of the population that is sampled determined exogenously to the agent. Define x_0 as the threshold value such that values of x equal or greater than x_0 are sampled: $x \geq x_0 = F_X^{-1}(1-\theta_0)$, where $F_X(x)$ is the cumulative distribution function of the marginal distribution of x: $F_X(x) = F_{X,Y}(x, \infty)$. The analyst knows the value of θ and sampled values of y, but not necessarily the values of x, so is generally convenient to express the expected value of the sampled value of the outcome in terms of θ rather than x: $E(y|\theta \leq \theta_0) = E(y|x \geq x_0)$.

2.2 Identifying the population mean

It is difficult to draw inferences about the population mean, μ, from $E(y|\theta \leq \theta_0)$ because a high value can be indicative of a high μ, but may also reflect an efficient sampling mechanism that identifies a higher share of high-y observations in a low-μ population. The marginal sampling efficiency at a specific sampling share (i.e. the expected value of the outcome for observations that are added to the sample when the sampling share increases infinitesimally) characterizes the nature of the purposive sampling and can be leveraged to infer information about the population mean (μ):

$$E(y|\theta_0) \equiv \lim_{\Delta\theta \to 0} E(y|\theta_0 \leq \theta \leq \theta_0 + \Delta\theta) = E(y|\theta \leq \theta_0) + \theta_0 \frac{dE(y|\theta \leq \theta_t)}{d\theta_t}$$

Since θ_0, $E(y|\theta \leq \theta_t)$ and $\frac{dE(y|\theta \leq \theta_t)}{d\theta_t}$ are all observable to the analyst, we can also treat $E(y|\theta_0)$ as observable. Exogenous variation in θ across subgroups that have the same joint distribution of $F(y,x)$ can be used to identify the marginal sampling efficiency $E(y|\theta_0)$, which can provide a plausible basis for drawing inferences about μ.

2.3 Statistical inference

Provided that we observe outcomes for subsamples of the population that are selected using the same sampling method, it is instructive to observe how $E(y|\theta \leq \theta_t)$ varies with changes in θ_t. If $\frac{dE(y|\theta \leq \theta_t)}{d\theta_t} = 0$ then this implies that $\mu = E(y|\theta \leq \theta_0)$. Intuitively, if an increase in the sample average of the outcome remains unchanged when the sampling share is increased, then this suggests that the sampling occurs in a haphazard way that is unable to effectively distinguish between high-y and low-y observations. Provided that all sampling shares between 0 and 1 are drawn with the same sample selection mechanism, then $E(y|\theta \leq \theta_0)$ provides an unbiased estimate of μ. On the other hand, if we can reject the hypothesis that $\frac{dE(y|\theta \leq \theta_t)}{d\theta_t} = 0$ in favor of $\frac{dE(y|\theta \leq \theta_t)}{d\theta_t} < 0$, then sampling can be concluded to occur based on informative values of x.

The simplest approach to point-identifying μ is to express the observable subgroup means in the general form as a function of θ_t rather than as a joint function $F_{X,Y}(x,y)$:

$$E(y|\theta \leq \theta_t) = \int_{-\infty}^{\infty} \frac{yP(\theta \leq \theta_t|y)f_Y(y)}{\theta_t} dy$$

and then to make distributional assumptions about $P(\theta \leq \theta_t|y)$ and $f_Y(y)$. When the outcome is binary, the conditional probability simplifies to:

$$P(y=1|\theta \leq \theta_t) = \frac{P(\theta \leq \theta_t|y=1)\mu}{\theta_t}$$

which only requires assumptions about the distribution of $P(\theta \leq \theta_t|y=1)$ in order to identify μ. If x is known to be partly determined by the value of the outcome, then

expressing x in error form as $x = \beta_0 + \beta_1 y + e$ where $E(e|y) = 0$ makes it possible to write the observable sample share of the outcome as:

$$P(y = 1|\theta \leq \theta_t) = \frac{P(e \geq F_X^{-1}(1 - \theta_t) - \beta_0 - \beta_1)\mu}{\theta_t}$$

A distributional assumption for the error term e is therefore all that is required to identify both $\beta \equiv \beta_0 + \beta_1$ and μ.

3. Context

TB has been the leading cause of death for over a decade in South Africa but the lack of reliable estimates of local TB prevalence makes it difficult to allocate government resources efficiently (Statistics South Africa 2011). This is especially important for MDR-TB which accounts for around 3% of all TB patients in South Africa, but consumes about 50% of the TB budget (WHO 2011; Pooran et al. 2012). MDR-TB treatment success rates in South Africa are close to 45% compared to 79% for drug susceptible TB (WHO 2014).

Conventional thinking about estimating MDR-TB prevalence focuses on two avenues that are expensive and logistically complex: increasing the frequency and coverage of TB prevalence studies and expanding access to high-technology testing (e.g. Xpert) so that everyone can be tested for drug resistance (Cohen et al. 2008, Weyer et al. 2013). Theron et al. (2015) calls for better use of existing data to inform tailored responses in the fight against TB, however advances in this area have been slow. For over a decade there have been calls for more MDR-TB prevalence studies yet only 8.3% of the population in the 27 high MDR-TB burden countries live in an area where at least two accurate population surveillance data points are available to estimate trends (Cohen et al. 2014).

Official guidelines counsel clinicians to screen patients for TB based on symptoms (current cough, weight loss, night sweats, fever) which are non-specific to TB (i.e. a noisy signal) and do not differ for patients with MDR-TB (Department of Health 2013). Before Xpert drug resistance testing was widely available in South Africa, the guidelines indicated that *only* those with the highest risk of MDR-TB should be tested, which included people with TB symptoms who had been previously treated for TB, patients who had failed TB treatment, and those who were known contacts of MDR-TB cases (Department of Health 2013). Clinicians can ascertain a patient's approximate risk of MDR-TB from a medical history and physical exam to determine a patient's MDR-TB risk relative to other patients before ordering laboratory testing. The risk factors are a good but imperfect signal of MDR-TB so many cases could initially go undetected. Previously, MDR-TB was primarily due to acquisition through incorrect or incomplete treatment, however recent evidence from the period of study shows that most incident MDR-TB cases are due to transmission instead, which makes risk factors worse predictors of MDR-TB positivity (Kendall et al. 2015). Resource shortages such as stockouts of test materials or drugs, long lab wait times or heavy clinical workload may prevent some at-risk patients from being tested for MDR-TB, thereby inducing exogenous variation in the proportion of TB patients who are tested for MDR-TB.

4. Methods

4.1 Data

We use data from South Africa's National Health Laboratory Service (NHLS) database on TB tests performed on patients aged 16-64 in public health facilities (hospitals and health clinics) for the period January 2004 - September 2010, which includes over 11 million patients. Our analysis sample comprises 2,190,780 TB-positive test records from patients in 5,122 health facilities (249,779 of which are tested for MDR-TB). For TB and drug susceptibility testing, the data include the type of test performed, test result, testing facility location, test date and basic patient demographics. We consider TB-positive results from culture testing and smear microscopy. Patient records are linked using unique patient identifiers created by the NHLS. Our dataset spans 8 years of frequent observations, which allows us to observe several sudden policy shifts that affected the inclination to test for MDR-TB.

In some analyses, we use data only from new patients in order to limit the sample to one diagnostic episode per patient, exclude treatment monitoring tests, and examine the sample without a history of TB testing (which in most cases implies without a history of TB treatment). Ethics approval was obtained from the University of Michigan Institutional Review Board and the University of Cape Town Faculty Ethics in Research Committee.

4.2 Identification

We apply our model to the clinician's decision to perform drug susceptibility testing on a patient with suspected TB. Clinicians search for MDR-TB-positive patients (i.e. high-value targets) because a confirmed diagnosis is required to initiate MDR-TB patients onto treatment. Prior to 2011, guidelines did not recommend testing all TB-positive patients for MDR-TB and MDR-TB testing resources were highly limited. In the absence of resources to test every patient for drug resistance, the clinician's testing decision is based on their knowledge of the patient's risk factor profile and the policy guidelines in place, subject to having time and resources available for drug resistance testing. Suppose the clinician observes information about the patient's underlying propensity to have MDR-TB in the form of a noisy signal (x) that includes non-specific symptoms such as cough and fever. The clinician can use this observed signal to determine a patient's relative likelihood of having MDR-TB conditional on being TB-positive.

The proportion of individuals who can be tested for MDR-TB, θ, is exogenously determined by institutional factors (such as in national testing guidelines, and the availability of test materials and human resources) which naturally has implications for the type of individual who gets tested, but should have no effect on the clinician's ability to rank patients or on the underlying share of drug-resistant patients. Clinicians will

triage patients and order drug resistance testing for the proportion θ of patients with the highest expected likelihood of having MDR-TB based on the noisy signal.

The assumption of exogenous changes in the proportion of patients tested over time implies a set of moment conditions that can be used to identify the parameters of interest. In this case we can use the period dummy variables as instrumental variables. However, we may be concerned that identifying off changes over time may reflect underlying transmission dynamics of the epidemic that drive the true prevalence. We therefore also use instruments based on exogenous institutional variation in MDR-TB testing rates. The national policy changes are orthogonal to facility-level variation in the lagged proportion of patients tested and lagged MDR-TB prevalence because their timing is neither determined by clinician decision making nor by deviations from the underlying prevalence trend. Intuitively, these instruments represent discontinuous changes in θ that cannot, in the short term, be correlated with relatively smooth trends in prevalence or the noise-to-signal ratio.

We include the following policy changes: MDR-TB surveillance study results reported (Jan 2002); national anti-retroviral therapy (ART) for AIDS rollout begins (July 2004); WHO declares TB an emergency in Africa (August 2005); first poster on XDR-TB presented at Conference on Retroviruses and Opportunistic Infections (CROI) (February 2006); South African government National Strategic Plan released (January 2007); clinical guidelines require one rather than three negative smears for smear-negative diagnosis (January 2009). We also use the facility-level and local-area-level availability of ART as instruments to incorporate sub-national exogenous variation. With the exception of the ART rollout, these policy changes are highly unlikely to change the composition of population of people present at health facilities or affect the ranking ability of clinicians and should therefore serve as valid instruments.

4.3 Estimation

In our empirical analysis we assume that *e* in the equation $x = \beta_0 + \beta_1 y + e$ a standard normal distribution *e*~nid(0,1) and we normalize β_0 to 0. This is sufficient to identify the model parameters signal-to-noise ratio $\beta = \frac{\beta_1}{\sigma_e}$ and population mean μ with observable sample variation in $P(y = 1|\theta \leq \theta_t)$ induced by different values of θ_t across time. We estimate the model first assuming that μ is constant over time, and then allowing μ to take a quadratic function of time where the prevalence is restricted to fall within the unit interval:

$$\mu_t = \Phi\left(\mu_0^* + \mu_1^* t + \mu_2^* t^2\right),$$

where $\Phi()$ represents the standard normal cumulative distribution function and *t* is the time period. Estimation of this model is complicated by the fact that x follows a mixed normal distribution and its inverse will not generally have a closed-form solution:

$$F_X(x) = \mu\Phi(x\text{-}\beta) + (1-\mu)\Phi(x).$$

The inverse, $F_X^{-1}(1-\theta_t)$, must therefore approximated using either simulations or numerical approximation techniques. We apply two approaches to estimate the model

parameters: a method of simulated moments estimator (in which F_X^{-1} is approximated using simulations) and a maximum likelihood estimator (in which a Hermite polynomial of degree 4 is used).

4.3.1 Method of simulated moments estimation

We can define the model error term (i.e. the difference between observed and predicted outcomes) as

$$u_{it} = y_{it}\text{-}P(y_{it} = 1|\theta \leq \theta_t).$$

The method of simulated moments (MSM) estimator exploits the exogeneity assumption:

$$E(u_{it}|\Omega) = 0$$

where Ω represents all the information available to the clinician at the time of the testing decision. In the sample data this implies that

$$E[z_{it}\{y_{it}\text{-}P(y_{it} = 1|\mu, \beta, \theta_t)\}] = 0$$

where z_{it} is a vector of instrumental variables, the elements of which are believed to be orthogonal to the individual's likelihood of having MDR-TB. If we define the GMM moment function as

$$g_{it}(y_{it}, z_{it}, \mu, \beta, \theta_t) = z_{it}\{y_{it}\text{-}P(y_{it} = 1|\mu, \beta, \theta_t)\}$$

then the generalized method of moments estimator can be expressed as

$$\arg\min_{\mu,\beta} \left(\sum_{t=1}^{T}\sum_{i=1}^{N} g_{it}(y_{it}, z_{it}, \mu, \beta, \theta_t)\right)' W \left(\sum_{t=1}^{T}\sum_{i=1}^{N} g_{it}(y_{it}, z_{it}, \mu, \beta, \theta_t)\right)$$

where W is the weighting matrix. We cannot calculate $P(y_{it} = 1|\mu, \beta, \theta_t)$ analytically, but replacing this probability with its simulated counterpart in the GMM estimator allows us to estimate the model parameters with the method of simulated moments. The estimates are obtained from 2,000 draws from a Halton sequence. The initial parameter values are obtained after performing a series of grid searches to obtain the most promising parameter space.

4.3.2 Simulated maximum likelihood estimation

Similarly, we estimate our model parameters using maximum likelihood estimation where the log-likelihood is expressed as

$$\log L(\mu, \beta)$$
$$= \sum_{i=1}^{N}\left[y_{it}\log\left(\frac{\mu}{\theta_t}\{1 - \Phi(F_X^{-1}(1 - \theta_t) - \beta)\}\right)\right.$$
$$\left. + (1 - y_{it})\log\left(1 - \frac{\mu}{\theta_t}\{1 - \Phi(F_X^{-1}(1 - \theta_t) - \beta)\}\right)\right]$$

In this model we use the MSM estimates as our initial parameter values and a Hermite polynomial of degree 4 to approximate the value of F_X^{-1}.

5. Results

Figure 1 shows that the proportion of TB-positive patients tested for MDR-TB (•) and the proportion that test positive among those tested (▲) are negatively correlated, both in long-run trends and short-term fluctuations. These patterns are consistent with our

theoretical model in which clinicians triage TB patients for testing based on the observed likelihood of being MDR-TB. An increase in the tested proportion (θ) implies extending the test to patients deemed less likely to have MDR-TB by the clinicians, in other words the signal-to-noise ratio $\beta>0$ and $\frac{dE(y|\theta\leq\theta_t)}{d\theta_t} < 0$.

Figure 1 also shows that the percent of all TB-positive patients who were diagnosed with MDR-TB (■) tracks the percent of all TB-positive patients who were tested for MDR-TB (•) reasonably well over this period. As more TB-positive patients are tested for MDR-TB, more MDR-TB cases are found (consistent with our model in that θ is a limiting factor and the signal-to-noise ratio $\beta<\infty$) and the share of MDR-TB tested patients who are MDR-TB-positive falls ($\beta>0$).

The percentage of all TB-positive patients (based on smear, culture or PCR) who were tested for MDR-TB (•) was fairly stable between 8-10% from 2004-2006, spiked up at the end of 2006 and again at the end of 2007 before steadily increasing from the end of 2008 to 2010, when it reached 23%. The percentage of all TB-positive patients who tested positive for MDR-TB (■) was stable between 1-1.3% from 2004-2006 and rose up to 2% in 2010. MDR-TB cases as a percentage of all those tested for MDR-TB (▲) was steady at around 12% until late 2007 when it rose to 15% and then steadily declined.

To further investigate the sampling mechanism, Figure 2 plots the sampling efficiency (solid line) as θ varies between approximately 8% and 23% in different subsamples. As expected, our results show that the prevalence in the selected sample falls when θ rises (i.e. $\frac{dE(y|\theta\leq\theta_t)}{d\theta_t} < 0$). The dashed line represents the marginal sampling efficiency – the expected value of y for observations that are added to the sample when θ increases infinitesimally. It declines up to approximately 18% of the TB+ patients being tested for MDR-TB, and then rises slightly at the highest values of θ.

We estimate that MDR-TB prevalence in South Africa was as high as 3.02-3.57% (Table 1) over this period, which is as much as 1 percentage point higher than the 2011 WHO estimate of 2.5% based on case notification rates (WHO 2011). This indicates that between 17 and 30 percent of all MDR-TB cases went undetected during this period. The standard errors for our estimates are small. When we relax the assumption that MDR-TB prevalence is constant over time, the MSM-IV estimates show that MDR-TB prevalence rose from 3.9% in 2004 to a peak of 4.3% from mid-2006 to mid-2007, and then fell to 3.6% by the end of 2010 (Figure 3, Table 1). The MSM estimates using time dummies and policy changes are very similar, indicating that changes in testing resources over time were fairly exogenous. MSL estimates were in general about 0.5 percentage points lower than the MSM estimates. The signal-to-noise ratio in the clinician-observed noisy signal of risk factors is estimated to range between 0.64-0.79. MSM-IV and MSL methods produce very similar estimates of MDR-TB prevalence (μ) and the signal-to-noise ratio (β) in the full sample.

Figure 4 provides evidence of the validity of our estimation method because the time pattern of MDR-TB prevalence predicted by our model matches the observed MDR-TB prevalence reasonably well in both the long and short run. This shows that the majority of the variation in MDR-TB prevalence can be explained by changes in the proportion of patients tested (θ) alone. The match is worse where the observed prevalence has more peaks and troughs (2007-2010).

New patients and patients with a previous test result have different underlying noise-to-signal ratio and estimated MDR-TB prevalence. The MSL results show a prevalence of 3.0% for new patients and 6.6% for repeat patients. Though the prevalence survey distinguishes between new and retreatment cases, which we are not able to do in our data, our results for new patients are still close to prevalence survey estimates for new cases (2.1%, 95% CI: 1.5%-2.7%) as are our results for repeat patients compared to retreatment cases (4.6% CI 95%: 3.2%-6.0%) (Centre for Tuberculosis 2016). As expected, the values for β reflect that clinicians have less information upon which to assess the risk profile of the new patients compared to repeat patients. For the MSL, β is estimated at 0.623 for new and 1.026 for repeat patients.

6. Discussion

Our results indicate that the assumptions about clinician behavior in our theoretical framework are consistent with the data. Figure 1 shows that clinicians do prioritize testing patients that are more likely to be MDR-TB positive, but that this prioritization is imperfect. Figure 2 shows, as expected, that the sampling efficiency falls as a greater proportion of the population is sampled. Our simple framework is able to match observed patterns in the data very closely (Figure 4). The fact that our MSM-IV-policy change estimates differ little from the MSM-IV time dummy estimates provides evidence to support our assumption that the constraint on diagnostic testing resources (θ) changes exogenously over time. The fact that our MSM-IV results change little with the addition of an instrument related to the rollout of ART, which occurred at the facility level and varied geographically and temporally, provides additional support that μ and β are well-identified. Our results do not exhibit characteristics indicative of weak instruments – large standard errors or sensitivity to changes in the sample – therefore concerns about bias in the estimates due to an influence of local MDR-TB prevalence on MDR-TB testing rates appear unfounded (Stock, Wright and Yogo 2002). This method can be further validated with applications to other data sources.

Our estimates of MDR-TB prevalence are between 17 and 30 percent higher (0.5-1 percentage point) than the WHO 2010 estimate of 2.5% (WHO 2011) which was calculated by adjusting the number of cases observed (i.e. notification rates) using case detection rates determined by "expert opinion". Because our data do not have full coverage of KwaZulu-Natal , which likely has the highest MDR-TB burden, our estimates are a lower bound on the true national MDR-TB prevalence. As expected, our results show that MDR-TB prevalence increased following the 2001-02 prevalence study, which found an MDR-TB prevalence of 2.9% overall and 6.6% in the population with a history of TB treatment (Weyer et al. 2007). Extrapolating our MSM-IV-policy change

time-varying prevalence results finds an MDR-TB prevalence of 3.4% for 2001-02 which falls within the CI of 2.4-3.5% from the prevalence survey, and 2.8% for 2012-2014 which exactly matches the estimate from the 2012-14 prevalence survey (Weyer et al. 2007, Centre for Tuberculosis 2016). In light of these results, it is likely that resource allocation to MDR-TB between 2002-2016 was sub-optimal due to under-estimates of MDR-TB prevalence. Additional resources should be therefore be allocated to the National Tuberculosis Program to increase efforts to control MDR-TB.

From a health policy perspective, high rates of under-detection of MDR-TB highlight the need for additional diagnostic resources and MDR-TB treatment for new cases that are identified. The current MDR-TB budget allocation is therefore likely to be insufficient. In addition, our new MDR-TB estimates should be used as input parameters for TB modeling studies that inform health policy because MDR-TB prevalence is often highly influential in these models (see Acuna-Villaorduna et al. 2008, Vassall et al. 2011, Meyer-Rath et al. 2012, Dowdy et al. 2014). Finally, more frequent prevalence surveys are needed to track the evolution of MDR-TB prevalence over time. Prevalence surveys and rigorous statistical analysis of routine data are complements rather than substitutes: recently completed MDR-TB prevalence studies can serve to further calibrate methods such as ours, and estimates from the analysis of routinely collected data can inform the design of prevalence studies to maximize precision and minimize cost.

4.1 Data Limitations

Our study population is the same as for the TB prevalence studies: individuals who present at a public health facility, are determined to be at risk for TB and have TB testing performed. Both will underestimate the prevalence of TB and MDR-TB in the population to the degree that cases do not present to health facilities, or are overlooked as at-risk by health workers, or due to diagnostic tests not being perfectly sensitive. Though the data have been deduplicated using an algorithm devised by the NHLS, poor patient linking across time may lead to double counting of MDR-TB patients and bias our estimates upwards. If clinicians order drug susceptibility testing only after treatment failure has been observed, then in the data clinicians will appear to have better information (stronger signal value) than they actually do. In the absence of prevalence study benchmarking, this would bias our estimates upwards.

4.2 Conclusion

This study developed a novel econometric method for estimating disease prevalence from routinely collected data. We found that approximately one-quarter of MDR-TB cases in South Africa were undiagnosed between 2004-2010 which contributed to high transmission rates and high TB mortality rates. These findings demonstrate the need for increased investment in early detection of MDR-TB, such as the ongoing implementation of Xpert technology, and more effective treatment, such as new antibiotics (WHO 2014).

Our results underscore the importance of continuous surveillance that accounts for under-detection rather than simply relying on notification rates in order to ensure that the health

system is diagnosing as many MDR-TB cases as possible. The fact that our method relies solely on existing routine data, which is widely available and inexpensive to collect, and can be run using standard statistical software makes it a particularly attractive way to produce estimates. In addition, the data requirements for our method are minimal compared to alternative prediction or imputation methods since patient characteristics need not be observed by the analyst. Notably, our method does not require that all patients or a random sample of patients be screened for the disease which makes it a lower-cost alternative to increasing the frequency of prevalence surveys.

This straightforward, yet powerful, approach to disease surveillance is simple and adaptable enough to be applied to many infectious and non-infectious diseases in the developing world where prevalence data is lacking. Our method can be applied to MDR-TB where access to Xpert drug resistance technology is limited, and to extensively drug resistant TB where Xpert is available but testing for resistance to additional first- and second-line drugs is less common. It is also a viable strategy for identifying localities and patient groups with high disease burdens, especially when prevalence survey samples are too small to produce precise local prevalence estimates. Routine statistical analysis results can also function as an early warning system for outbreaks, especially if they are able to discern deviations from the prevalence trends over time. Countries such as India, China and Russia, which together account for over 50% of MDR-TB cases worldwide could benefit from the deployment of this method to target investments in MDR-TB diagnosis, such as the rollout of new diagnostic technologies or new second-line TB drugs.

Statistical analysis of routinely collected data is an inexpensive and effective way to monitor the prevalence of any number of population characteristics. Our approach is widely applicable because of the minimal data requirements, wide availability of routinely collected data and abundance of policy changes to serve as valid instruments. Our straightforward, yet powerful, approach can be used to evaluate the effectiveness not only of clinicians, but also of tax authorities, customs officials and law enforcement. Ultimately, using routinely collected data to monitor population prevalence and agent effectiveness is a high-value strategy to guide evidence-based policy making and implementation in resource-limited settings.

Acknowledgements

We thank Sue Candy, Ananta Nanoo, Michelle Potgeiter and Andrew Whitelaw for assistance with the data and helpful comments. We thank Michael Budros, Elizabeth Brouwer, David Ederer, Kathryn Fischer, Alana Sharp and Jifang Zhou for research assistance. We thank Jacob Bor, Florian Marx, Tyler McCormick, Tom Moultrie, Edward Norton, Kate Schnippel, Margaret Triyana, Sean Wasserman, Mark Wilson, seminar audiences at the University of Michigan and Stellenbosch University, and participants at the Union for African Population Studies conference, University of KwaZulu Natal Microeconomic Analysis of South Africa conference and European Workshop on Econometrics and Health Economics for helpful feedback. Financial

support was provided by the University of Michigan School of Public Health Global Public Health Program, the Center for Global Health and the University of Michigan Health Management and Policy Department McNerney Award.

References

Acuna-Villaorduna, C., Vassall, A., Henostroza, G., Seas, C., Guerra, H., Vasquez, L., ... & Gotuzzo, E. (2008). Cost-effectiveness analysis of introduction of rapid, alternative methods to identify multidrug-resistant tuberculosis in middle-income countries. Clinical Infectious Diseases, 47(4), 487-495.

Bärnighausen, T., Bor, J., Wandira-Kazibwe, S., & Canning, D. (2011). Correcting HIV prevalence estimates for survey nonparticipation using Heckman-type selection models. Epidemiology, 22(1), 27-35.

Cameron, A.C, and P.K. Trivedi. Microeconometrics: methods and applications. Cambridge university press, 2005.

Centre for Tuberculosis, 2016, South African Tuberculosis Drug Resistance Survey 2012–14, Johannesburg: Centre for Tuberculosis.

Churchyard, G. J., L. D. Mametja, L. Mvusi, N. Ndjeka, A. C. Hesseling, A. Reid, S. Babatunde, and Y. Pillay. (2014) "Tuberculosis control in South Africa: Successes, challenges and recommendations." SAMJ: South African Medical Journal 104(3): 234-248.

Clark, SJ., and Houle B., "Validation, Replication, and Sensitivity Testing of Heckman-Type Selection Models to Adjust Estimates of HIV Prevalence" PLoS ONE 9 (11) e112563. doi: 10.1371/journal.pone.0112563(2015)

Cohen, T, et al. "Challenges in estimating the total burden of drug-resistant tuberculosis." American journal of respiratory and critical care medicine 177.12 (2008): 1302-1306.

Cohen, T, et al. "On the spread and control of MDR-TB epidemics: An examination of trends in anti-tuberculosis drug resistance surveillance data." Drug Resistance Updates 17.4 (2014): 105-123.

Department of Health of South Africa, January 2013, Management of Drug Resistant Tuberculosis Policy Guidelines.

Dowdy, D. W., Houben, R., Cohen, T., Pai, M., Cobelens, F., Vassall, A., ... & White, R. (2014). Impact and cost-effectiveness of current and future tuberculosis diagnostics: the contribution of modelling. The International Journal of Tuberculosis and Lung Disease, 18(9), 1012-1018.

Greene, W. H. "Econometric Analysis." 6th ed. (2008).

Hogan, DR et al. "National HIV prevalence estimates for sub-Saharan Africa: controlling selection bias with Heckman-type selection models" Sexually Transmitted Infection 91.8 (2015) i17-i23

Judd, Kenneth L. Numerical methods in economics. MIT press, 1998.

Karim, S S A, Churchyard G J, Karim Q A, Lawn S D. HIV infection and tuberculosis in South Africa: an urgent need to escalate the public health response. Lancet 2009; 374(9693):921-933.

Kendall, E. A., Fofana, M. O., & Dowdy, D. W. (2015). Burden of transmitted multidrug resistance in epidemics of tuberculosis: a transmission modelling analysis. *The Lancet Respiratory Medicine*, 3(12), 963-972.

Kim, J.Y., A. Shakow, K. Mate, C. Vanderwarker, et al. 2005. "Limited Good and Limited Vision: Multidrug-Resistant Tuberculosis and Global Health Policy," Social Science and Medicine Journal, 61: 847–859.

Klopper, M et al. "Emergence and Spread of Extensively and Totally Drug-Resistant Tuberculosis, South Africa" Emerging Infectious Diseases 19.3 (2013) 449-455

Manski & Lerman 1981. Structural Analysis of Discrete Data with Econometric Applications, Editor D. McFadden, Cambridge: M. I. T. Press, 1981.

McFadden, D. 1989. "A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration," Econometrica, 57(5): 995-1026.

McGovern, ME et al "On the Assumption of Bivariate Normality in Selection Models *A Copula Approach Applied to Estimating HIV Prevalence"* Epidemiology 26.2 (2015) 229-237

Medecins Sans Frontieres, Partners in Health and Treatment Action Group. An evaluation of drug-resistant TB treatment scale-up. July 2011.

Meyer-Rath G, Schnippel K, Long L, MacLeod W, Sanne I, et al. (2012) The Impact and Cost of Scaling up GeneXpert MTB/RIF in South Africa. PLoS ONE 7(5): e36966. doi:10.1371/journal.pone.0036966

Nyirenda M, et al. "Adjusting HIV Prevalence for Survey Non-Response Using Mortality Rates: An Application of the Method Using Surveillance Data from Rural South Africa" PLoS ONE 5 (8): e12370,doi: 10.1371/journal/pone.0012370

Phelps, E. S. (1972). The statistical theory of racism and sexism. American Economic Review, 62(4), 659-661.

Pooran A, Pieterson E, Davids M, Theron G, Dheda K (2013) What is the Cost of Diagnosis and Management of Drug Resistant Tuberculosis in South Africa? PLoS ONE 8(1): e54587. doi:10.1371/journal.pone.0054587.

Sakarovitch, C et al. "Estimating incidence of HIV infection in childbearing age African women using serial prevalence data from antenatal clinics" Statistics in Medicine in Wiley InterScience (2007) 320-335

Statistics South Africa. Mortality and causes of death in South Africa, 2010: findings from death notification. Statistical Release P0309.3. http://www.statssa.gov.za/publications/ p03093/p030932010.pdf Pretoria, South Africa: Statistics South Africa, 2013. Accessed October 2014.

Stock, J. H., Wright, J. H., & Yogo, M. (2002). A survey of weak instruments and weak identification in generalized method of moments. Journal of Business & Economic Statistics 20(4).

Theron, G et al. How to eliminate tuberculosis 1 "Data for action: collection and use of local data to end tuberculosis" The Lancet 386.10010 (2015) 2324-2333.

Vassall A, van Kampen S, Sohn H, Michael JS, John KR, et al. (2011) Rapid Diagnosis of Tuberculosis with the Xpert MTB/RIF Assay in High Burden Countries: A Cost-Effectiveness Analysis. PLoS Med 8(11): e1001120. doi:10.1371/journal.pmed.1001120

Weyer, K, et al. "Rapid molecular TB diagnosis: evidence, policy making and global implementation of Xpert MTB/RIF." European Respiratory Journal 42.1 (2013): 252-271.

Weyer, K., Brand, J., Lancaster, J., Levin, J., & Van der Walt, M., 2007, "Determinants of multidrug-resistant tuberculosis in South Africa: results from a national survey." South African Medical Journal 97(11).

World Health Organization, 2011, WHO Report 2011 Global Tuberculosis Control, Geneva: World Health Organization.

World Health Organization, 2013, Multidrug-resistant tuberculosis (MDR-TB), 2013 Update, March 2013.

World Health Organization, 2014, WHO Report 2013 Global Tuberculosis Control, Geneva: World Health Organization.

Table 1: Estimated MDR-TB prevalence (μ) and signal-to-noise ratio (β).

	Signal to noise ratio (β)	MDR TB rate (μ)	Prevalence index coefficient (μ_0^*)	Time coefficient (μ_1^*)	Time squared coefficient (μ_2^*)
MSM-IV: policy changes					
No trend	1.071***	0.0302***			
Quadratic trend	0.765***		-3.201***	0.017***	-0.001***
	(0.0185)		(0.0450)	(0.0049)	(0.0002)
MSM-IV: time dummies					
No trend	0.889***	0.0357***			
	(0.0092)	(0.0003)			
Quadratic trend	0.793***		-3.231***	0.017***	-0.001***
	(0.0084)		(0.0297)	(0.0036)	(0.0001)
MSL					
No trend	1.071***	0.031***			
	(0.0076)	(0.0076)			
Quadratic trend	0.640***		-3.02***	0.019***	-0.001***
	(0.0848)		(0.1083)	(0.0068)	(0.0003)

*Notes: Table presents coefficients and standard errors. Sample includes TB-positive patients ages 16-64 in public health facilities from January 2004-September 2010. N = 2,190,780. *** - Significant at the 1% level, ** - 5% level, * - 10% level.*

Table 2: Estimated MDR-TB prevalence (μ) and noise to signal ratio (β) using MSL.

	Repeat patients		New patients	
	Signal to noise ratio (β)	MDR TB rate (μ)	Signal to noise ratio (β)	MDR TB rate (μ)
No trend	1.026***	0.066***	0.623	0.030**
	(0.0377)	(0.0074)	(0.4325)	(0.0146)

*Notes: Table presents coefficients and standard errors. Sample includes TB-positive patients ages 16-64 in public health facilities from January 2004-September 2010. N = 2,190,780. New patients defined as within three months of first TB test in data. *** - Significant at the 1% level, ** - 5% level, * - 10% level.*

Figure 1: Percent of TB-positive cases tested for MDR-TB, percent of TB-positive cases MDR-TB-positive, and percent of MDR-TB-tested cases MDR-TB-positive from National Health Laboratory Service data (scaled to two Y-axes to show how the testing rate and testing-positive rate track reasonably well over time).

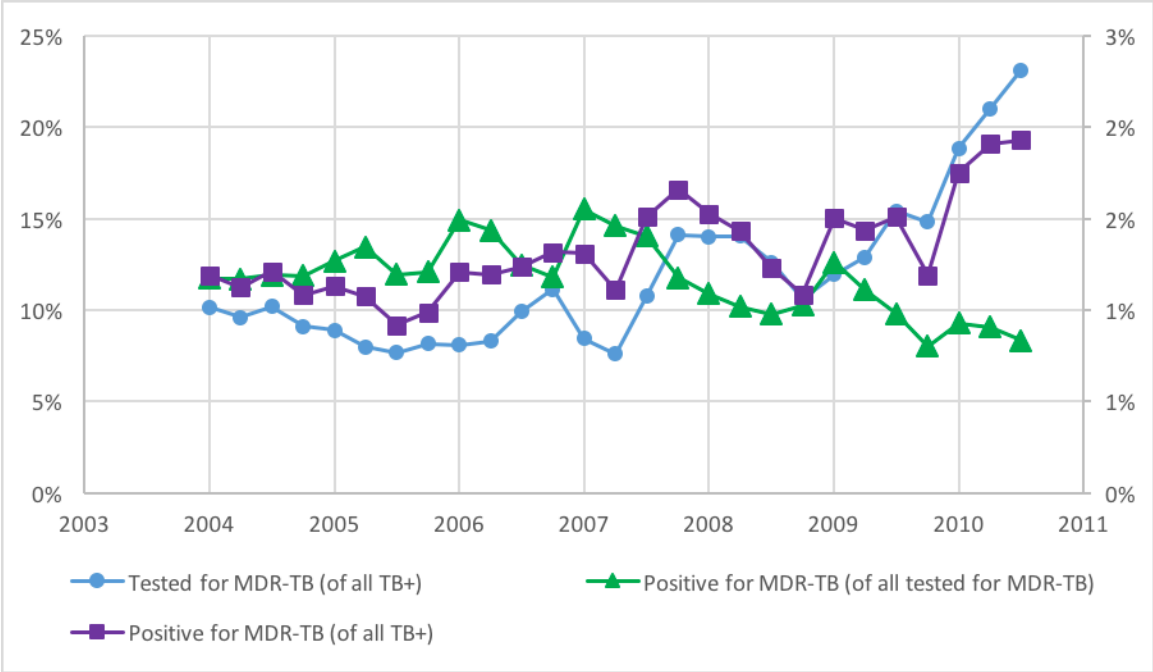

Figure 2: Sampling efficiency and marginal sampling efficienct at different proportions of TB+ patients being tested for MDR-TB (θ).

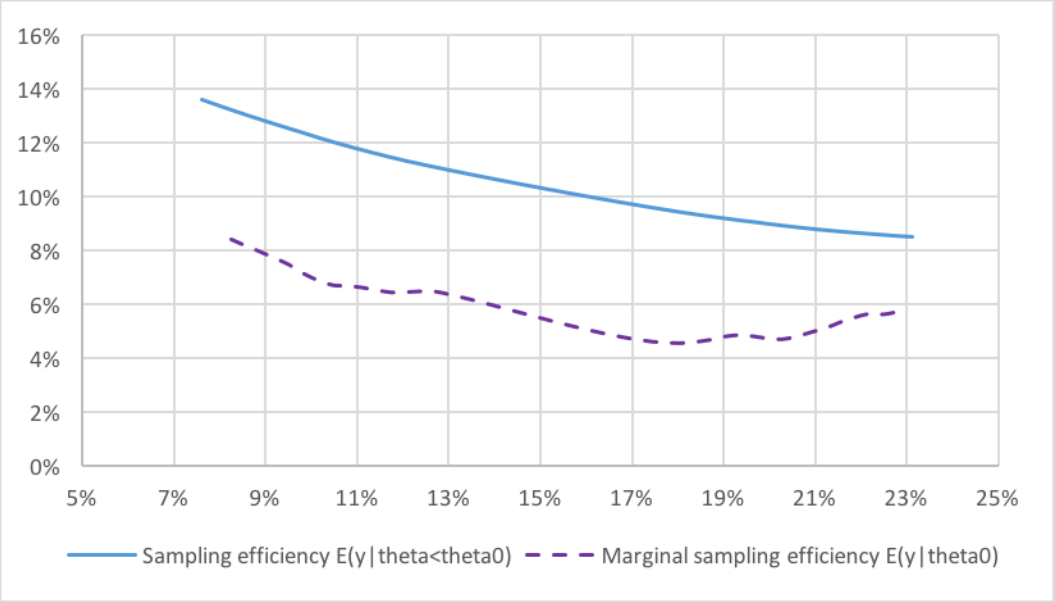

Figure 3: Predicted time trends in MDR-TB prevalence (%) in South Africa estimated from NHLS data using method of simulated moments and simulated maximum likelihood estimation.

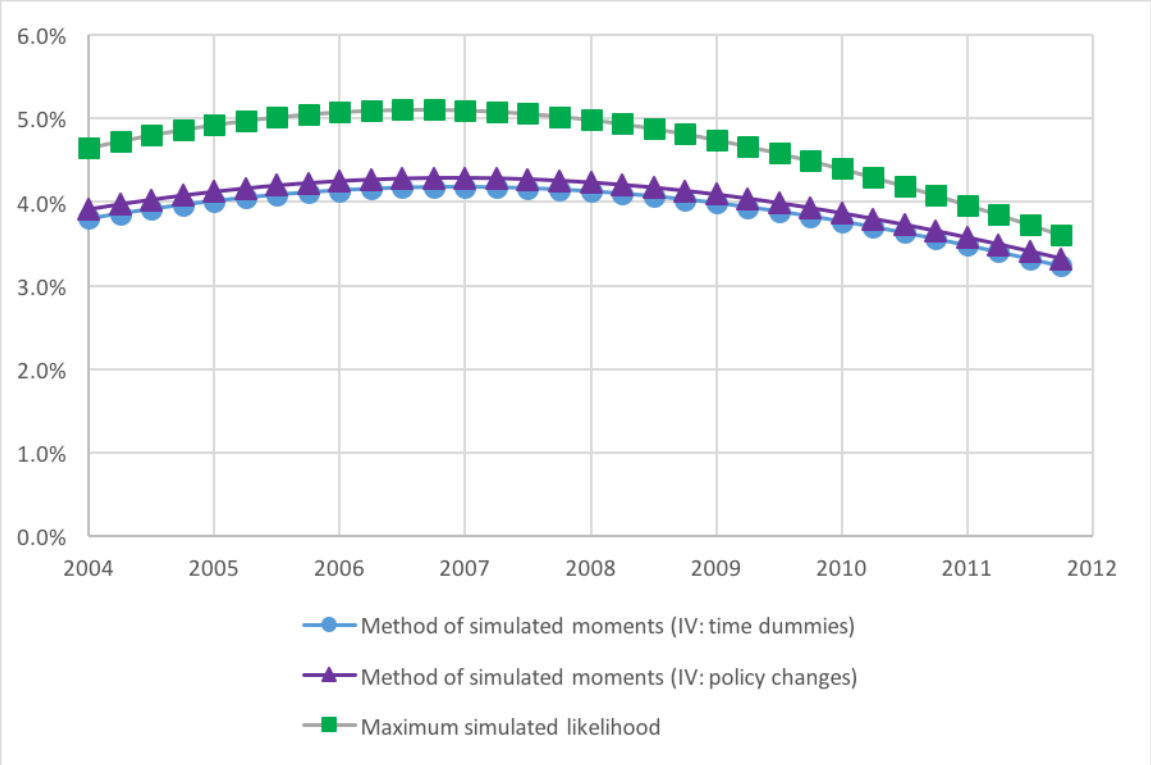

Figure 4: Observed MDR-TB prevalence over time in NHLS data compared to MDR-TB prevalence estimated from simulated maximum likelihood estimation.