Mortality
Felipe Menares

The beginning of the world's demographic transition occurred in northwest Europe, where mortality began a secular decline around 1800. In many low-income countries of the world, the decline in mortality began in the early twentieth century and then accelerated dramatically after World War II. The historical decline can be summarized as the result of knowledge, science and technology; specifically, poor countries benefited from health innovations developed in richer countries. On the contrary, there are views stating that the growth in population from the late 1700s was not because of advancements in the field of medicine or public health, but instead due to improvements in overall standards of living, especially diet and nutritional status. This is known as the Mckeown thesis, which states that improvements in the standard of living, nutrition and wellbeing, rather than public health or medicine-oriented policies which impacted the health of populations.

The main discussion is shaped around the relationship between health, income and mortality. While it is clear that there is an effect of the first two factors on mortality, it remains an open question whether and how health and income relate to each other. *Does health cause higher income or is it income that causes better health?* This question has been studied by scholars that have addressed it both between countries and within countries. Between countries, the concave relationship between life expectancy and national income is well established. It is possible to observe that among poor countries, increases in income are associated with increases in life expectancy, but as income rises, the relationship levels off. There are also shifts in this relationship suggesting that changes in longevity are related to exogenous factors rather than solely income. For instance, people in low-income countries, for a given increment in income, live substantially longer today than they did before.

Nowadays, when most of technology has reached the developing world, the relationship between income and health becomes crucial to mortality inequalities when studied *within* countries, rather than in between countries. While there is a vast literature studying the effect of income inequality on population health, there are no direct links. Therefore, within countries the people with lower earnings who are sick, explain much of the correlation between income and health, rather than a causal relationship from higher income to better health. Correlations come from factors other than income inequality itself, some of which are linked to broader notions of inequality and inequity that are most likely important for health.

There is evidence using administrative data that in the United States, higher income was associated with greater longevity, and differences in life expectancy across income groups increased over time, being correlated with health behaviors and local area characteristics. In that sense scholars have stressed underlying characteristics to address mortality disparities. For instance, (i) the use of education instead of income as the mechanism to explain inequalities, (ii) racial disparities through the study of the opioid crisis, (iii) place, such as the hispanic paradox and the place of birth of migrants.

Likewise, others find that disparities were caused primarily by a number of chronic diseases and injuries with well-established risk factors, concluding that disparities in life expectancy cannot be

explained by race, income, or basic health-care access and utilization alone. As a matter of fact, beyond medical advances, the major factor in reducing cardiovascular disease mortality is the reduction in smoking. Therefore, the most revealing underlying cause of health inequalities is smoking, explaining why we see some of the well-established mortality gradients. The term is used to emphasize that there are “graded” differences in health running across ranked groups, not just between poor and rich.

Health care improvements are being well studied to address how this affects health and consequently mortality. Many of the studies claiming to show a causal effect of health insurance on health do not do so convincingly because the observed correlation between insurance and good health may be driven by unobservable factors. Using the language of econometrics, the problem is that health insurance is an endogenous variable. Therefore, it is necessary to address this question by understanding the difference between observational, natural (quasi-experiments) and experimental studies.

Most of the recent work done is based on the Affordable Care Act reform (ACA). In 2014, the ACA expanded eligibility for the Medicaid program to include all adults in families with incomes under 138 percent of the Federal Poverty Level. Previously only pregnant women, adults with disabilities, and very low income parents tended to qualify for Medicaid coverage. Although intended to apply to all states, a 2012 Supreme Court decision made the Medicaid eligibility expansion optional. As a result, only 29 states and the District of Columbia expanded coverage in 2014, with 7 additional states electing to expand over the next several years. Despite non-universal adoption, approximately 13.6 million adults gained Medicaid coverage.

These studies take advantage of the variation in state adoption of this large expansion in coverage to compare changes in mortality among individuals in expansion states and non-expansion states. Therefore, most of the studies answered these questions with quasi-experimental design, by implementing event study and differences-in-differences analysis, as well as the use of propensity score. While there is a variety of findings among the results, broadly, they find either (i) no effect or a (ii) positive effect. No effects are mostly due to sample problems and the consequential lack of statistical power, while the positive effect results rely mostly on the elimination of the pre-trends in order to claim causal effects.

Latin America is not the exception, there are well developed observational studies showing spatial and economic disparities as well as the study of health insurance on mortality. However, most of the reforms are based on expansion, rather than explicitly aiming towards the reduction in mortality. Therefore, to the best of my knowledge there is not yet a clear relationship between health care reforms and mortality inequalities. I believe Chilean Health Insurance reform is a good setup to answer this question. This was based on a universal health plan which provides explicit healthcare guarantees with regard to coverage for 85 health conditions with high mortality. This is not a health reform on the coverage aspect but rather in the treatment of diseases.

Mortality
Felipe Menares

Showpiece: **Did the ACA Medicaid Expansion Save Lives? Borgschulte & Vogler (2020)**

The authors estimate the effect of the Affordable Care Act (ACA) expansion on county-level mortality during the first four years following the reform. Medicaid expansion removes categorical exclusions and bases eligibility solely on income at or below 138 percent of the federal poverty level. The authors use restricted-access microdata covering all deaths in the United States, which is paired with population denominators by county, age group, sex, and race from the U.S. Census and then merged with county-specific economic and demographic variables, such as the unemployment rate, the poverty rate, and real median income. Even though the unit of study is the county, the identifying variation comes from the state level timing of adoption.

At first sight, there is a pre-trend in all-cause and amenable-cause mortality rates in expansion counties relative to non-expansion counties. They investigate this issue in more detail with specific causes of deaths and conclude that the confounding pre-trends are largely explained by differences in cardio-vascular and respiratory mortality. To address these pre-existing differences, they use event study and difference-in-differences (DD) models in conjunction with propensity score reweighting, to select a group of control counties in a way that balances pre-treatment characteristics, including mortality outcomes.

The event-study model enables them to both test for the presence of pre-trends and capture the evolution of the treatment effect over time. The pooled DD results averages the year-specific effects estimates in the event study into a single contrast of pre-and post-expansion differences between treatment and control counties. The coefficient of interest represents the change in mortality among counties in states that expanded Medicaid in 2014 relative to counties in states that did not. The term county-TOT should not be confused with the individual-TOT that can be calculated if given the size of the affected population.

The timing of the effect aligns with the expansion and shows mortality reductions became larger over time. They find a reduction in all-cause mortality in ages 20 to 64 equaling 11.36 deaths per 100,000 individuals, a 3.6 percent decrease. This estimate is largely driven by reductions in mortality in counties with higher pre-expansion uninsured rates and amenable causes of death, likely to be influenced by access to healthcare. They also cannot reject the null that the coefficients for amenable and non-amenable mortality are statistically different.

Using a first-stage estimate of a 4.15 percent decrease in uninsurance in the sample and adjusting for the higher mortality of the Medicaid population implies that the expansion prevented 30 percent of deaths in this population, a large reduction in mortality. Taken together, these calculations imply that one death was averted for every 310 newly-covered individuals 20 to 64 years of age and 180 newly-covered individuals aged 55 to 64 years. They said that it is important to note that access to healthcare may have larger effects over a longer time and numerous ways in which it affects a larger population. They find evidence consistent with effects

accruing in a broader population than those who receive insurance as a result of Medicaid expansion.

They performed several other exercises for robustness: First they discuss results from a wide range of alternative specifications, such as, matching procedures, definition of the treatment, county trends and trimming the propensity score. Second, they conduct a permutation test to investigate the likelihood that the model is identifying the causal effects or whether the effects found are simply due to chance as a result of potentially confounding issues such as model misspecification or lack of power. Third, the importance of leverage and outliers. Finally, estimates using elderly adults as an additional control group, in a DD and DDD on age 65-75, similar to BHNS. None of the results differ from the results found in their preferred specifications.

A cost-benefit analysis shows that the improvement in welfare due to mortality responses may offset the entire net-of-transfers expenditure associated with the expansion. They calculate the benefits of expansion under two sets of assumptions: first, by applying point estimate for all-cause mortality to the ages 20 to 64; and second, under a conservative scenario in which their results apply only to amenable cause mortality for the 55-to-64 age group.

Comments

The authors address one of the most important questions nowadays: the effects of health insurance in preventing mortality. However, in this context, the main challenge is to identify causal effects, therefore in the absence of an experimental design where there is random assignment to treatment - i.e. medical insurance- they address this question by using a quasi-experimental design. Specifically, they use the Affordable Care Act, and its unequal implementation in 2014, as the major reform that had heterogenous timing of adoption which can be exploited. In order to do that, they performed an event study and a DD model combined with propensity score reweighting.

The identification assumption for causality under a quasi-experimental DD design is the existence of flat pre-trends pre-reform between the treated and untreated, namely counties in states that adopted the ACA and those that didn't. However, there are clear pre-trends in all-cause and amenable-cause mortality rates in expansion counties relative to non-expansion counties in the years prior to 2014. Therefore, they select a group of control counties in a way that balances pre-treatment characteristics, including mortality outcomes, based on economic, demographic, and political characteristics. When observing this issue in more detail with specific causes of deaths, the confounding pre-trends are explained by differences in cardiovascular and respiratory mortality.

The main challenge of the paper is to have a good control group through a propensity score matching, otherwise the use of solely an event study or DD approach would be sufficient. They tackle this problem with state-of-the-art techniques to choose the covariates explaining these

differences. The authors claim that these features of the implementation of the research design serve to discipline the analysis and add credibility to the findings. They clearly stated the difference between the results found, saying that the term county-TOT should not be confused with the individual-TOT (ITOT) that can be calculated if given the size of the affected population, adding credibility to the results.

However, all their results depend on how well they are explaining what makes treated and non treated counties different. As a result, their coefficients will have additional noise, because in order to explain the flat pre-trends, they need to implement a propensity score, deviating the attention to the matching procedures rather than the actual quasi-experimental design. Consequently, the propensity-score step will not necessarily improve the across-county comparison, in terms of both bias and efficiency. In particular, counties with propensity scores that are outside the overlap region are trimmed from the sample. Therefore, the trimming performed, which they said follows recent literature, complicates the extrapolation of these results nationwide.

They also compare this paper with some of the recent literature relative to the quasi-experiments. On one hand, it's important to note that while OHIE is the large scale randomization study, it lacks sufficient statistical power to detect mortality effects of the magnitude documented by the authors, since the ACA expansion was nationwide and affected a much larger share of the population. On the other hand, compared to Black et al 2019 (BHNS here after), they use the double-lasso method to specify the propensity score model. This data-driven model-selection procedure prevents overfitting, thereby increasing power relative to research designs assessed in BHNS.

The empirical strategy and data used differ from the models evaluated in the power calculations in BHNS in three key ways. First, it focuses on a data-driven propensity-score model compared to models with evident pre-trends, which, they acknowledge, significantly reduces power. Second, it uses two additional years of data relative to the power calculations in the most recent version of BHNS. Given the growing effects, this increases the likelihood of detecting an effect both due to an increase in the precision and likely magnitude of the effects. Finally, they report reduced-form estimates for uninsurance and mortality while remaining agnostic as to their ability to estimate precise treatment-on-the-treated estimates. In contrast, BHNS draw their strongest conclusions (especially in the introduction of their paper) when using a first-stage estimate that is likely considerably smaller than the group of individuals actually affected by the reform.

This paper’s contribution to the discussion is clearly remarkable, albeit their heavy reliance on the propensity score, the use of machine learning in order to avoid overfitting, bias and improve efficiency, adds credibility to their results. They also provide sufficient evidence to believe that their results are good enough to claim a causal positivity effect of health insurance on mortality, performing a vast amount of robustness analyses to check their implementation. Alternatively, it would be useful to compare their results with additional robustness analyses, such as, synthetic control and geographical regression discontinuity for problematic areas, such as the south.