Restricting Economic Activity to Slow the Growth of Covid-19

A Cost-Benefit Analysis

By James Sly

June 12th, 2020

Adapted from Economic Policy During and After a Pandemic (Sly 2020)

Introduction

The first economic problem presented to decisionmakers during a pandemic is figuring out how much to restrict economic activity in order to make critical gains in public health. For the most part, these decisions have been left to governors based on advice from public health experts, since these are the experts who have the best insights and knowledge into how to achieve improvements in public health. This is exactly as it should be, however, any decision on social distancing restrictions also has an effect on the economy as well as public health, which is within the purview of economists. Figuring out how to weigh the tradeoffs is fundamentally a simple cost-benefit analysis, which is also within the skill set of economists.

Three Ways the Pandemic Can End

This policy memo attempts to lay out the basic costs and benefits of imposing certain restrictions on the economy. In any pandemic, there are basically three ways it can end. First, very strict limits on economic behavior can be enacted that reduce the transmission rate to the point that the number of new cases shrinks dramatically over a period of time. These restriction can continue until the number of new cases actually reaches zero, or until the number of new cases is low enough that they can be dealt with through a rigorous system of testing, contact tracing, and isolation.(1) Second, more moderate limits on economic behavior can be enacted that cause the number of new cases to plateau over time, and then this moderate level of infection continues until a vaccine is eventually developed and administered. (2) Third, the virus can be allowed to spread undeterred throughout society in order to achieve herd immunity more quickly, at which point the virus has trouble finding new people to infect, and the number of new cases shrinks down to zero once the infection rate gets high enough.(3)

Estimating the Value of an Individual Human LIfe

The basic tradeoff between these three options weighs the potential benefits of lives saved against the financial cost of additional economic restrictions.(4) In order to be able to do a basic cost benefit analysis, you need to set an overall value for an individual human life, so that the public health benefits can be converted into a financial value that can then be compared to the financial costs from enacting more restrictions on the economy. In general, the federal government uses a value for an individual human life ranging from $7.9 million to $9.6 million depending on the department, (5) which for our purposes we will round up to $10 million as an upper bound for our estimate.

That statistic estimates the value of a whole human life, but it is possible to argue that most of the people who die from Covid-19 are relatively old, so perhaps we should develop an estimate for each

year of life saved. If the average person lives approximately 80 years in the United States (6) and a total life is worth $10 million, then each year is worth approximately $125,000. When deciding which treatments to cover, Medicare uses a value of about $50,000 for each extra year of life, but one study suggests it should be more like $129,000. (7) This cost-benefit analysis is attempting to be simple, so for the purposes of this paper, we will use $100,000 per year of life as the estimate for the lower bound.

Then we need to figure out how many years are lost with each Covid-19 death. For this, we can use data coming from New York City that says about 50% of deaths are from those age 75 or older, about 25% of deaths are from age 65-74, and about 25% of deaths are from those under age 65. (8) Assuming each person would have lived to about 80 years old, this means the average loss of life due to a death from Covid-19 is about 10 years. (9) Once you do the math, that means each death from Covid-19 probably leads to a loss of life worth about $1 million at the lower bound, and $10 million at the upper bound.

How Many Lives Can Be Saved by Restricting Economic Activity?

In the next step of our cost-benefit analysis, we need to figure out how many lives would be saved with each level of economic restriction. Imposing the strictest economic rules, as they did in China, generally keeps the number of deaths to a minimum, where in China less than 5,000 people died, while the major European countries were less successful. Their approach to social distancing restrictions resulted in about 33,000 deaths in Italy, 28,500 deaths in France, 27,000 in Spain, while Germany had less than 10,000 deaths. Scaled up to the population of the US, this means our country could have seen the number of deaths range from about 1,000 to 200,000. (10) The second option, which is currently being used in the US, maintains moderate restrictions that cause the number of cases and deaths to plateau. This option has resulted in about 100,000 deaths over the last two months, which if this rate lasts for a total of 12 months until a vaccine is developed, this would result in about 600,000 deaths in total.(11) The third option, where the US tries to get herd immunity as quickly as possible, could see infection rates rise to around 40-70% of the population or about 130-230 million people.(12) If the death rate ranges from about 0.66% to 1.15% for each infected person, then the total number of deaths resulting from the herd immunity approach ranges from about 1 million to 2.5 million.(13)

Comparing the Most Restrictive Strategy to the Least Restrictive Strategy

If we were to compare option #1, the most restrictive, to #3, the least restrictive, we know the least restrictive has minimal economic cost due to social distancing, and now we need to estimate the economic cost of the most restrictive option. In the most hard-hit city in China, Wuhan, the lockdown lasted about two and a half months, and Italy has partially relaxed its own national lockdown after about two and a half months as well.(14) We can round that up to around 3 months and basically just guess (since there is little data available yet) that the lockdown reduced economic output by about 20%. (15) If this impact were applied to the US economy this would mean the response to the virus would reduce annual economic output by about 5% over an entire year, costing the country about $1 trillion in a $20 trillion dollar economy.

We know from before that option #1, the most restrictive strategy, leads to somewhere between 1,000 and 200,000 deaths, which has a public health cost ranging from $1 billion on the low end to $2 trillion on the high end using our low and high estimates for the value of a human life. The economic cost of the most restrictive option is approximately $1 trillion so the total cost would likely range from $1 trillion to $3 trillion.(16) The economic cost from option #3 is minimal, because it is the least restrictive,

but the public health cost is substantial, ranging from $1 trillion to $25 trillion based on the estimate of losing 1 million to 2.5 million lives.(17) That means the total cost for the most restrictive option is approximately the same as the least restrictive option under the most cautious assumptions, and represents less than one-eighth of the cost under the most generous assumptions.(18) In this case, the cost-benefit analysis points to the idea that the most restrictive option is preferable to the least restrictive option for nearly the entire range of estimates.

Comparing the Most Restrictive Strategy to the Moderately Restrictive Strategy

The next step in our cost-benefit analysis is to compare option #1, the most restrictive, to option #2, the one with moderate restrictions. We already know the economic and public health costs for option #1. For the economic costs for option #2, we will just have to guess that this will reduce economic output by about 10% for an entire year until a vaccine becomes available, which leads to an economic impact worth around $2 trillion.(19) The public health cost of 600,000 additional deaths represents anywhere from $600 billion to $6 trillion, which means the total cost ranges from $2.6 trillion to $8 trillion. (20) The total cost for option #1 ranged from $1 trillion to $3 trillion, so clearly the most restrictive option will be better than the option with moderate restrictions since both the economic shutdown is less costly (because it is shorter) and has a smaller public health impact as well by saving at least 400,000 lives.

Comparing the Moderately Restrictive Strategy to the Least Restrictive Strategy

Even if the most restrictive option might be the least costly among all three options, political pressures may prevent our own democracy from imposing such harsh limits on our own economic and social freedoms. That means option #1 might not be possible for political reasons. For this reason, we also need to compare option #2, the one with moderate restrictions, to option #3, the least restrictive option. We already calculated the total cost of option #2, which ranged from $2.6 trillion to $8 trillion, and we already calculated the total cost of option #3, which ranged from $1 trillion to $25 trillion. Using the high estimate for the death toll of the least restrictive strategy, option #3, means the moderately restrictive strategy, option #2, is preferred when the value of a human life rises above $1 million. Using the low estimate for the death toll for option #3, means that the moderately restrictive strategy, option #2, is preferred when the value of a human life rises above $5 million.(21)

This means that even if option #2 is preferred over a wider range of estimates, we can still do even better by looking a little closer at what this means. If we assume that the value of an individual life is perfectly chosen to make both options equally costly, then under these circumstances the average value of those restrictions in public health terms is equal to the cost of the economic damage. Being an economist of course, we are not interested in the average costs and benefits, we are interested in the marginal costs and benefits. If we divide up option #2 into a number of smaller restrictions, some of those restrictions will have the public health benefits exceed the economic costs, but some of the will have the economic costs outweigh the public health benefits, and once you add them all up at once, then they turn out to be exactly equal. That means if the value of a human life is set to make the options exactly equal, then that means the most valuable restrictions are probably worth it, but some of the other ones are not, so when looking at the marginal cost and benefits, it might make sense to do only half the restrictions. That way you reduce output by only 5% of GDP over the course of a year rather than 10%. These estimates are not accurate enough to determine if you should only do

restrictions that impose an economic cost up to 2% or up to 8%, but perhaps at low estimates of a value of human life, the restrictions should impose an economic cost less than 10%, and so 5% would seem like a reasonable guess.

If, however, we use the high estimates for the value of a human life, then the public health benefits of the moderately restrictive option might be two to nine times the economic costs,(22) at which point we might want to impose additional restrictions that cause an economic cost above 10%, say up to 15%. Again, the estimates are not accurate enough to determine a precise number, but it is safe to say that if the value of a human life is relatively low, than there will be some restrictions that make economic sense, but perhaps only those that cause a 5% reduction in output, but if we use a high estimate for the value of human life, then perhaps we should impose restrictions that cause a 15% reduction in output.

Conclusion

Based on the cost-benefit analysis, it then becomes clear that the most restrictive option is probably the least costly of the three, since there are tremendous public health benefits and the shorter duration of those restrictions reduces the overall economic costs. If, however, the most restrictive option is not possible for political reasons, then a moderate restriction is preferred over a broader range of estimates, but once you look at the marginal cost of each individual restriction, policymakers might only want to implement restrictions that impose costs as low as 5% of output over a year, or as high as 15% of output over a year. In that case, the best you can do is evaluate each specific restriction to see if it provides a large public health benefit or a small public health benefit compared to the economic damage imposed, and try and mitigate the public health costs as much as you can while still maintaining a significant benefit for the economy.

End Notes

#1 – A variety of different countries have been able to reduce the number of new infections down to extremely low levels, where this particular long term outcome has been reached by Iceland, New Zealand, Taiwan, Vietnam, Luxembourg, Thailand, and China, who had fewer than 10 new infection on May 25, 2020, along with Greece, Switzerland, Australia, Austria, Norway, South Korea, Israel, Finland, Ireland, Denmark, and Japan, who all had fewer than 50 new infections on the same date based on data from worldometers.info.

#2 – The United States and United Kingdom are the most prominent countries that have seen the number of new cases plateau over time, where they have seen slow declines once the dramatic exponential rise ended in early April, but have not been able to achieve dramatic reductions in the number of new cases each day even by the end of May.

#3 – Even though Sweden has not officially announced they are adopting a strategy of herd immunity, they have been noticeably reluctant to enact social distancing restrictions that would be able to slow the spread of the disease, relying on voluntary measures instead. On May 25th, 2020, Sweden had 597 new cases of Covid-19, compared to less than 50 for the other Nordic countries like Denmark, Finland, Norway, and Iceland.

#4 – Even though there are other public health costs besides fatalities as a result of the pandemic, the impact of hospitalizations and intensive care admissions are substantially lower than dying, where the cost of each individual death is likely in the millions of dollars, as opposed to a cost in the tens of thousands of dollars for hospitalizations or ICU admissions. Since the cumulative hospitalization rate, according to the CDC, was around 67 per 100,000 people since March 1st, while the death rate in the US, according to worldometers.info was around 30 per 100,000 people, the cost of each hospitalization is only 0.1% to 10% of an individual death, and occurs only a bit more than twice as often as a fatality.

#5 – According to Wikipedia, the FDA uses a value of human life of $7.9 million when doing their cost benefit analyses, while the EPA uses $9.1 million. In 2014, the Department of Transportation used a value of $9.2 million, while in 2016, this rose to $9.6 million.

#6 – According to the CDC, life expectancy in the United States was 78.6 in 2017, which I rounded up to approximately 80 in order to make the calculations simpler.

#7 – Medicare uses a value of $50,000 for each additional quality adjusted life year (QALY) when deciding to approve new medical treatments for coverage, but recently Stanford professor Stefanos Zenios argued that this value should go up to $129,000 by looking at the life expectancy benefits of kidney dialysis.

#8 – As of May 13, 2020, the New York Department of Health reported that 48.7% of deaths in New York City were age 75 and above, 24.9% between the ages of 65 and 74, 22.4% between the ages of 45 and 64, and 3.9% between the ages of 18 and 44 and 0.06% below age 18 according to worldometers.info.

#9 – Life expectancy at birth was 78.6 in 2017 according to the CDC. Life expectancy at age 75 was 11.7 years in 2007, while life expectancy at age 65 was 18.6 in the same year. I assumed that the life expectancy for all people who died over the age of 75 was 5 years, that the life expectancy for those dying between the ages of 65 and 74 was 10 years, and those dying at age 64 or under was approximately 20 years. This is perhaps 5 years below the actuarial estimates of life expectancy at those particular ages, but this is designed to be a simple calculation, and many of those who died had underlying health conditions that would suggest a shorter life expectancy than the actuarial tables for the whole population. If the number of years lost was actually 15 rather than 10, it means the financial estimate of the cost of an individual death based on years of life would be $1.5 million, rather than $1 million, but this result could return to $1 million if an extra year of life was valued at $67,000, which is also a reasonable estimate. Using a higher value for the extra year of life would not fundamentally change the results, except the more restrictive options would look relatively more attractive when using the low estimate, because of the higher value put on the public health benefits they provide.

#10 – The data on total deaths and death rates per million in population from China and Europe due to Covid-19 come from the website worldometer.info as of March 25th, 2020. China had a death rate of 3 per 1 million people, which if scaled up to the US population estimates of 328 million from the US Census Department, then the US would have seen approximately 1,000 deaths. Spain had the highest death rate of 580 per 1 million people among the European countries mentioned, and if this were scaled up to the US population, then the US would have seen approximately 190,000 deaths. These numbers could rise since the pandemic is not yet over in Europe, but this is our best approximation based on the latest available data.

#11 – According to worldometers.info, deaths in the US first rose above 100 per day on March 22nd, 2020 and by May 25th, 2020 the number of deaths was 100,572, so approximately 100,000 people died in the United States over the preceding two month period. WebMD cites a variety of sources, including the CDC, which indicate that a vaccine will take at least 12 to 18 months to develop. If vaccine development began when the outbreak first took off in China in late January of 2020, then a vaccine might become available anytime between February and July of 2021. I assumed that a vaccine will become available 12 months after the virus took off in the US in late March, which is 14 months after vaccine development began and well within the range of possible outcomes. I also assume that deaths will continue occurring at about 50,000 per month for the entire 12 months until the vaccine arrives. The number of deaths has been trending down slightly over time in the US, but there is also the possibility of a second wave next winter, so this is just a general approximation of what might actually happen.

#12 – Harvard epidemiologist Marc Lipsitch estimated in a March 2nd interview with CBS News that 40% to 70% of the world's population could get infected with Covid-19, which if applied to the US population of 328 million, would lead to about 131 million to 230 million infections.

#13 – There are a variety of estimates of the infection fatality rate (IFR) for Covid-19. I looked at four estimates of the IFR, each one from a different area of the world. One was a calculation for New York State done by worldometers.info that estimated an IFR of 1.4%. Another was one done for Spain by the Carlos III Institute for Health which calculated an IFR of 1.15%. A third was published in The Lancet, Infectious Disease which calculated an IFR of 0.66% for China. A fourth study from the University of Bonn looked at coronavirus infections in Germany and calculate an IFR of 0.37%. I tried to make things simple by using a range that dropped the two most extreme estimates, and generated a range using the less extreme estimates. If you apply the exact percentages to the exact estimates of the number of infections based on the Marc Lipsitch proportions, you get the number of deaths ranging from 866,000 to 2.64 million which I rounded to 1 million to 2.5 million in order to make it simple and to also avoid a false sense of precision.

#14 - According to Wikipedia, the lockdown in Wuhan, China began on January 23rd, 2020 and was lifted on April 8th, 2020, a span of 77 days or about 2 ½ months. The nationwide lockdown in Italy started on March 9th, 2020 and was partially lifted on May 18th, 2020 (most businesses) with plans to end more restrictions on May 25th (swimming pools and gyms), and June 15th (theaters and cinemas). This means the lockdown lasted about 71 days until the partial relaxation started and at least 98 days until theaters are open, a span of approximately 2 ½ to 3 months.

#15 – In China, GDP in the first quarter declined by 6.8%, which if the lockdown only occurred during the last two months of the quarter, then GDP may have declined by about 10% during the lockdown period. This is likely a bit low, in part because Chinese statistics are not always accurate, and in part because some regions were hit worse than others, so the economic impact of the strictest lockdowns might be higher. The Economist predicts a decline in Italy's GDP of 7% for all of 2020, which would represent a decline of 28% if the entire decline happened over 3 months, but Italy's economy is likely not going to bounce back to previous levels immediately after the lockdown is over, so the loss in output during the lockdown is somewhat less than that. I guessed a drop in output of about 20% during the lockdown, which is somewhere in between the quarterly estimate for China and the yearly estimate for Italy after it is converted into a quarterly effect.

#16 – The low end of the estimate for the public health costs of most restrictive strategy, option #1, is to multiply the low death toll times the low value of life which is just 1,000 deaths times $1 million dollars, which adds up to $1 billion overall. The high end of the estimate for the public health costs of option #1 uses the high death toll and the high value of life, which multiplies 200,000 deaths times $10 million dollars, which adds up to $2 trillion. The low end of the total cost simply adds the economic cost of $1 trillion plus $1 billion for the low end of the public health cost, which gets you $1.001 trillion in total. The high end of the total costs simply adds the $1 trillion economic cost plus $2 trillion, which gets you $3 trillion.

#17 – The low end of the range of estimates for the total cost of the least restrictive strategy, option #3, is generated by taking the low end of the death toll, 1 million, and multiplying it by the low end of the value of life, $1 million, which ends up getting you a total of $1 trillion, since there is no economic cost attributed to option #3. The high end of the range of estimates for the total cost of option #3 is generated by taking the high end of the death toll, 2.5 million, and multiplying it by the high end of the value of life, $10 million, which ends up getting you a total of $25 trillion overall, since there is no economic cost included for this option.

#18 - This means the total cost at the low end of the estimates is $1.01 trillion for option #1 and $1 trillion for option #3, which means both are approximately the same. The total cost at the high end of the estimates is $3 trillion for option #1 and $25 trillion for option #3, which means option #3 is more than 8 times as costly.

#19 – The unemployment rate in the US rose from 3.5% in February 2020 to 14.7% in April of 2020 an increase of 11.2 percentage points. Retail sales in April fell 16.4% from the previous month and was down 21.6% from April of the previous year. This probably overstates the economic decline due to the pandemic since retail sales only accounts for half of about personal consumption, which represents only 70% of GDP, and other sectors of the economy are not as likely to be as hard hit. In comparison, the Congressional Budget Office recently projected that GDP would decline by 12% in the second quarter. A 10% guess for economic decline due to social distancing restrictions is generally in line with what CBO projects, rounded down to make the calculation simpler, which also reflects the great deal of uncertainty in the underlying estimate. To understand my assumptions about vaccine development please see end note #11. A 10% decline in output in a $20 trillion dollar economy leads to a total economic cost of about $2 trillion.

#20 – The range of estimates for the public health costs of the strategy adopting moderate restrictions simply multiplies the 600,000 estimated deaths by the low estimate of the value of human life, $1 million, to get the low estimate cost of $600 billion, and by the high estimate of the value of human life, $10 million, to get the high estimated cost of $6 trillion. The range of total costs simply adds the $2 trillion economic cost to the range of estimates for the public health costs.

#21 – Calculating the thresholds for the value of an individual life where option #2 becomes preferable to option #3 requires some simple algebra. When looking at the high estimate of the death toll for the most restrictive strategy, option #3, this option kills about 1.9 million more people (2.5 million minus 600,000) but saves about $2 trillion in economic costs, so these two values are equal when the value of a human life is worth approximately $1.05 million. When looking at the low estimate of the death toll for option #3, then in this case, option #3 kills about 400,000 more people (1 million minus 600,000), but

saves about $2 trillion in economic costs, so these two values are equal when the value of a human life is worth approximately $5 million.

#22 – For the moderately restrictive strategy, option #2, the economic costs are around $2 trillion, but the public health benefits are considerable. Option #2 saves about 400,000 to 1.9 million lives compared to the least restrictive strategy, option #3. If the value of a human life were worth $10 million, then the public health benefits would range from $4 trillion to $19 trillion, which is about 2 to 9.5 times the value of the economic costs of option #2, which is only $2 trillion.

References

Sly, James. 2020. "Economic Policy During and After a Pandemic." Unpublished Working Paper, June.

Worldometers.info. 2020. www.worldometers.info/coronavirus/#countries [Accessed May and June 2020]