

From Here to Full Inoculation: How an Epidemiological-Economic Model Can Help as We Roll Out Vaccines¹

Jeremy Kronick²

Paul Jenkins³

February 2021


In the early days of the COVID-19 global pandemic, with little information, the focus was rightly on health and on imposing a lockdown for all non-essential parts of the economy to minimize infection and mortality rates. While successful in turning infection rates back down and “flattening” the curve, it also resulted in the worst economic downturn since the Great Depression.

Success on the infection front was short lived, as rates began to surge again in the early fall months. Although the progress being made in the development of a vaccine is most encouraging, the rollout is likely to take much of 2021 – if not longer due to new variants. For at least another six to nine months, therefore, Canadians and policymakers must not let their guard down but remain vigilant in taking whatever precautions and measures are necessary to contain the virus and minimize deaths.

With the knowledge and experience gained since Canada went into lockdown in March 2020 – and, indeed, to some degree, with what was known even before then (see, for example, Ciuriak 2020) – we have reached a point where we can better gauge risks and have more informed discussions on the optimal mix of policies – on the health front, on the economic front and, critically, on the intersection of the two. Yet, deciding on the optimal mix of interventions, and communicating them in such a way as to earn and hold the trust of Canadians, remains a difficult task. Approaches have also varied considerably from one province to another.

The research presented in this Working Paper, using an epidemiological-economic model, is one tool that can be used to guide us through the period until a vaccine is widely available. By using data collected during the nine months between March and December 2020, we can use this model to ask what mix of policies would lead to health and economic outcomes such that one could not be improved without sacrificing some of the other – in economic parlance, a Pareto-optimal outcome (see Acemoglu et al. 2020, figure 1, on which our model is based). We do that by asking what intervention strategies, both by age and industry, would minimize economic loss without worsening, and even improving upon, the health outcomes Canada achieved before the latest infection surge.

Figure 1: The Policy Frontier



Source: Acemoglu et al. 2020, figure 1.

We find that, compared to a uniform lockdown in which everyone is locked down the same, a targeted approach to lockdowns by age or industry over the period until a vaccine is widely available would achieve the same or better health outcomes at lower economic cost; see Figure 1. A severe but short period of uniform, significant lockdown, as has been implemented in many areas of the country, would be possible even under a targeted strategy. New mutations would not change the improved results under a targeted approach, as we tested the impacts of higher transmission rates, which would capture new variants. Moreover, we find that interventions, such as improved testing and tracing, and mask wearing could further reduce the potential for spread, lessening the severity of necessary lockdowns and lowering economic costs without sacrificing lives.⁴

We begin by describing the model in terms of methodology and calibration, the calculation of economic loss under different economic (i.e., lockdowns) and non-economic

¹ Preliminary and incomplete. Do not quote without permission.

² Associate Director, Research, C.D. Howe Institute; jkronick@cdhowe.org

³ Former Senior Deputy Governor at the Bank of Canada, and Senior Fellow, C.D. Howe Institute; pjenkins@jenp.ca

⁴ For readers so inclined, one-page summaries of the main results in this paper can be found in Jenkins and Kronick (2020a, 2020b). We also note that the same theory holds when we reverse the question and ask what lockdown strategy would make the most sense if the economic losses were limited to X percent. Under a uniform approach, the health outcomes would be worse, since the economy would have to remain open

interventions (e.g., mask-wearing), and its limitations. We then present the model results along with variations, followed by a discussion of our robustness tests. Finally, we offer some insights for policymakers from the model results, along with some conclusions we feel are applicable to today’s situation.

Methodology

Our methodology is based on the work of Acemoglu et al. (2020), which, in turn, is based on the epidemiological SIR model literature.⁵ The key innovation in their model is the ability to treat age groups differently under different government intervention strategies (including full lockdown). Our model does two things differently: first, we calibrate the story for Canada; second, in addition to evaluating how targeted policies by age group affect the results, we look at how targeting policies by industry affect the results.

In our description here, we focus on the key elements of the Acemoglu et al. model, referring readers to their work for a more granular description.

Individuals can be divided into different risk groups (e.g., by age, by industry, etc.) $N_j = 1, \dots, J$, where j is, for example, young, middle aged and old, and N is the total population, which we normalize to unity – that is, $\sum_j N_j = 1$. At each point in time along our continuum, individuals may be susceptible (S), be infected (I), have recovered (R) or have died (D). At all times, then, the following holds:

$$S_j(t) + I_j(t) + R_j(t) + D_j(t) = N_j.$$

To go from susceptible to infected, one must obviously come in contact with an infected individual. An infected individual may or may not need ICU care. For our purposes, we label v_j as the percentage of people in each group who get infected and need ICU care. We assume this fraction differs by group but is constant within each group across time. Once in hospital, individuals either recover or die, while those who do not need hospitalization recover. Those outside hospital recover with a Poisson arrival rate γ_j , which is constant over time. So, for example, if $\gamma_j = 1/18$, then a person will recover on average in 18 days. Those in ICU recover with Poisson arrival $\delta_j^r(t)$, while those who die do so with Poisson arrival $\delta_j^d(t)$, both of which can vary over time. We assume the following is true:

$$\gamma_j = \delta_j^d(t) + \delta_j^r(t).$$

The implication here is that the proportions of those who need ICU care and those who do not do not change over time. Since the percentage of people who need ICU care is v_j , the total number of individuals in group j who need ICU care at any moment in time can be

enough not to exceed the economic loss constraint. A more targeted approach could strategically protect those most vulnerable.

⁵ The SIR model, first proposed by Kermack, McKendrick, and Walker (1927), models an infectious disease epidemic in a large population that is divided into compartments: “S” for susceptible, “I” for infected and “R” for recovered.

represented as $H_j(t) = v_j I_j(t)$, and so total hospital needs will equal $H(t) = \sum_j H_j(t)$. Since the probability of death is determined in part by these hospital needs, we write $\delta_j^d(t) = \psi_j(H(t))$ for some non-decreasing function ψ_j .

An important part of the risk of infection has to do with our ability to test, trace and isolate the infected. The better we do this, the lower the likelihood of transmission. We assume that, for those who do not require serious, ICU hospitalization, τ_j is the constant percentage of those individuals in group j we are able to detect and isolate. For those who do require ICU care, we set the constant detection and isolation percentage as ϕ_j . We also assume full recovery and immunity after infection.⁶ We thus define the percentage of those who fail to be detected and isolated as

$$\eta_j = 1 - (v_j \phi_j + (1 - v_j) \tau_j).$$

We set κ_j as the constant fraction of individuals who have been infected, have recovered and can now work normally.

We also introduce a couple of final health-related parameters: β , which gives us the reproductive rate of the disease absent any lockdown/intervention strategies, and ρ , which allows us to specify the contact rate between individuals. Where there is interaction between individuals in the same group, we write ρ_{jj} ; where individuals from one group interact with those from a different group, we write ρ_{jk} .

With these health elements in mind, we now highlight the economic impact of lockdowns or other intervention strategies. We assume that individuals in each group are able to produce w_j when the economy is open; under lockdown/intervention strategies, however, they are able to produce only $\zeta_j w_j$, where ζ_j is some fraction between 0 and 1. This implies that, under lockdown/intervention strategies, the productivity loss at home is $(1 - \zeta_j)w_j$ for each group.

To determine how much each group actually loses, we need to multiply this productivity loss, $(1 - \zeta_j)w_j$, by the lockdown/intervention strategy, which we label $L_j(t)$ and which can vary anywhere from 0 (no lockdown) to 1 (complete lockdown). For example, a full lockdown, $L_j(t) = 1$, would imply the entire productivity loss, $(1 - \zeta_j)w_j$, is incurred. We note briefly that no degree of lockdown can be fully effective, so even a full lockdown strategy would reduce interactions only by $1 - \theta_j L_j(t)$, where θ_j is our measure of inefficiency. It is this variable, $L_j(t)$, that is the choice variable for policymakers.

Lastly, we assume that a vaccine becomes fully available at the end of our simulation period, which we put at one year.⁷ At this point, for simplicity, we assume that a cure comes

⁶ We acknowledge that lingering side effects post-infection are possible, which would hurt productivity, and there have been documented cases of re-infections, bringing into question the length of immunity. For our purposes, however, over the one-year period we are looking at, it is sufficient to assume neither possibility would change markedly the results of the model.

⁷ The model does not incorporate an inoculation period. One can think of the availability a year out as incorporating both its availability and administration.

as well. Therefore, at the end of our simulation period, all those who were infected and susceptible have officially recovered.

Laws of Motion

With both the health and economic parameters in mind, we now explore how those susceptible, infected, recovered and died evolve over time in this model. In other words, what are the laws of motion for these variables (note that the dot above the left-hand side variable indicates a differential equation – that is, $\delta Y/\delta t$, or how Y changes when time changes by one unit, where Y is one of I , S , R , or D).

$$\dot{I}_j = \beta(1 - \theta_j L_j) S_j \sum_k \rho_{jk} \eta_k (1 - \theta_k L_k) I_k - \gamma_j I_j;$$

$$\dot{S}_j = -\dot{I}_j - \gamma_j I_j;$$

$$\dot{R}_j = \delta_j^r(t) H_j + \gamma_j (I_j - H_j); \text{ and}$$

$$\dot{D}_j = \delta_j^d(t) H_j.$$

A more complete discussion of the theory behind these laws of motion can be found in Acemoglu et al. (2020), but here we touch on a couple important points. With respect to the infection rate, in addition to the lockdown strategy the government employs, the transmission rate (which appears through β), the interaction rate (ρ) and our ability to detect and trace (η) are critical for keeping the growth of this number low. The evolution in the numbers of those who recover from infection and those who do not is influenced by the evolution of their respective arrival rates, which themselves are influenced by the pressure on the hospital system – meaning that there is a direct correlation to the rate of growth of infection in society. Hospital capacity itself, which influences the effect that growing infection rates have on recovery and death rates, is controlled for by the calibration we give to v_j , as discussed earlier.⁸

Economic Losses

We have described so far a health system in which the choice of lockdown leads to a specific evolution of those susceptible, infected, recovered or deceased among each of our different groups j . The health objective that will guide the policymaker in our model is the desire to minimize total lives lost, which we describe as

$$LivesLost = \sum_j D_j(T),$$

⁸ We acknowledge that not everyone dies in hospital – in Canada and elsewhere, COVID-19 infected elderly have been dying in long-term care homes – but it is beyond the scope of this model to deal with where people die.

where T is the final period, at which point the vaccine and cure arrive.⁹

Of course, there are economic losses to this story as well. At first, when we had minimal data on the disease, it was right, in our opinion, to be as cautious as possible by locking down non-essential activities and dealing with the ensuing economic losses. But now that we understand the disease and its spread considerably better (though by no means perfectly), we can introduce the economic loss component in the decisionmaking process and ask ourselves what policies might achieve the same or better health outcomes¹⁰ than occurred over the nine months between March and December 2020 but at a lower economic cost.

To answer this question, we need an economic loss function that a policymaker will try to minimize. For each group, we can define economic loss as follows:

$$\Psi_j(t) = (1 - \varsigma_j)w_j S_j(t)L_j(t) + (1 - \varsigma_j)w_j I_j(t) \left(1 - \eta_k (1 - L_j(t))\right) + (1 - \varsigma_j)w_j (1 - \kappa_j)R_j(t)L_j(t) + w_j \Delta_j v_j \delta_j^d(t)I_j(t).$$

This is a long, complicated formula, which is easier to understand if we break it down into pieces. The first term tells us how much loss there is from working from home for those who are susceptible. The second term does something similar for the infected, although these individuals also would create an economic loss simply from not being able to work due to being sick. The third term is 0, as long as all those who were infected and have recovered are identified and can go back to work – that is, $\kappa_j = 1$. Otherwise they are in the same predicament as those who are susceptible and experience the loss from working at home. The last term is the loss from those who die, which is a product of the death rate ($v_j \delta_j^d(t)I_j(t)$) in time t and the total loss of economic activity from dying over what the individual would have contributed over his or her remaining work years ($w_j \Delta_j$).

The policymaker will consider this economic loss function across groups over the course of the entire period before a vaccine and cure arrive in order to try to minimize

$$EconLoss = \int_0^T \sum_j \Psi_j(t) dt.$$

Lockdowns

For the policymaker, this is an optimal control problem over the period from time 0 to time T , which we set to be one year. In other words, each period, the policymaker must choose a lockdown strategy, $L_j(t)$, for the different groups j . Acemoglu et al. (2020) compare

⁹ Limiting infections instead of deaths is another approach that has some validity in that there are concerns over the long-term consequences of being infected with COVID-19. However, we do not have sufficient information on these consequences to model the optimal policy mix.

¹⁰ As a baseline, we focus on the “same” health outcomes achieved over the nine months from March to December 2020 in terms of total deaths as a percentage of the total population.

uniform policies – that is, all groups are treated the same ($L_j(t) = L(t)$) – versus targeted policies, where lockdowns take place differently among different groups.¹¹ Recall that a targeted approach does not preclude the possibility of everyone being severely locked down for a period of time, as we have seen in parts of the country.

We follow Acemoglu et al. (2020) by differentiating first by age group – young (20–49), middle (50–64) and old (65+) – and then by industry grouping – low, medium and high contact (see Appendix A for a breakdown of how we categorize each industry). To determine which industry fits in which category, we calculate, using Statistics Canada data, the percentage of people within each industry who were able to work from home prior to the pandemic. Industries where the percentage was between 0 and 10 percent were considered high contact, those between 10 and 20 percent were considered medium contact and those above 20 percent were considered low contact.

Our focus is on minimizing economic losses subject to maintaining the same or better health outcomes – measured by lives lost – to those Canada experienced between March and December 2020.¹²

Lockdowns versus Non-economic Interventions

The structure of the model is such that the policymaker takes as given a series of parameters and their calibration, and asks what is the optimal economic intervention strategy – in other words, how much should each age group/industry type be locked down. The policymaker has two options: either a uniform lockdown where every individual/industry is treated the same, or targeted lockdowns, in which case different intervention strategies across groups are possible. However, the model also allows for a comparison of economic losses under different parameter calibrations, enabling us to distinguish losses under different combinations of economic and non-economic interventions.

Economic interventions, as discussed above, include both uniform and targeted lockdowns by age group and industry; we note that, even under a targeted approach, all groups could be locked down severely and equally for a period of time. The non-economic interventions we focus on include testing and tracing, and transmission rates via mask wearing. These are “non-economic” in the sense that they can improve health outcomes without negatively affecting the economy. The model permits evaluating these interventions by altering the value of η , which represents the failure to detect and isolate an infected individual, and β , which is the key parameter for determining the transmission rate.

¹¹ Our focus, which Acemoglu et al. (2020) discuss, is semi-targeted lockdowns in which those older than age 65 are treated differently from the young and middle aged; the latter two groups are treated the same. We do this given the clear separation it allows between working-age individuals and those eligible for retirement.

¹² Like Acemoglu et al. (2020), we use a non-linear dynamic programming method called IPOPT, using APMonitor. For more information, see Hedengren et al. (2014) and Wachter and Biegler (2006).

Calibration

To run our dynamic optimization model, we needed to calibrate the key parameters listed above. These parameters remain the same under both the uniform and targeted lockdown approaches. As noted just above, some parameters can be adapted in scenarios where we want to compare different non-economic interventions. Appendix B has the complete list of these parameter values and an explanation of their sources.

Model Limitations

These models can be used to simulate a wide range of possible interventions that can return infection rates to the low levels seen in summer 2020 (or better), while minimizing economic losses. That said, any such modelling exercise is a simplified description of a process of cause and effect, with a wide confidence band around its parameters.

Recognizing these limitations, our objective with the results presented below is not to be prescriptive, but to help identify how combinations of interventions undertaken in a strategic way could turn the current situation around and guide us to better health and economic outcomes as Canada moves along the path toward full vaccine rollout.

Bearing this in mind, the results below should be seen not as projections, but as notional, in that they give direction and orders of magnitude of economic losses under different intervention strategies. Put differently, the “art” in model development and use is knowing what issues any given model is designed to address and then using the model in such a way as to extract the right insights.

Model Results

From peak to trough – the end of February to the end of April 2020 – Canada’s gross domestic product (GDP) contracted by 17.7 percent, or \$354 billion, at an annualized rate (or \$88.5 billion at a quarterly rate). It subsequently rebounded by 10 percent in the third quarter. In terms of the level of GDP, however, the economy is still below pre-COVID levels.

To compare our model results to Canada’s actual experience requires some care.

As discussed above, the essence of our results is in comparing economic losses under different intervention strategies, with each strategy achieving the same health outcomes. With that in mind, the data presented below are best thought of as estimates of economic losses peak to trough. In that sense, our results can be compared to the actual economic losses Canada experienced before the economy began to rebound in late spring 2020.

At the same time, our results are also intended to provide insights on future policy. In other words, the starting conditions of our simulations are designed to reflect the resurgence in infection rates during fall 2020. With our one-year simulation horizon, the model results therefore can be interpreted as a comparison of economic losses under different intervention strategies to get us from where we are to the point of widespread inoculation.

Finally, the baseline results in this Working Paper are for Canada as a whole. As the current situation varies considerably across provinces and territories, and even within each jurisdiction, we also investigate the different current conditions and approaches of two large provinces with different testing and tracing capabilities – namely, Ontario and British Columbia – in order to compare their estimated economic losses under different strategies.

Age-based Results

We start with the age-based results. We view these results as our primary display of the effects of targeting, as we have seen from the beginning of this crisis a clear distinction of the disease’s impact on different age groups. Moreover, the industry results we present in the next section naturally use data only for working-age individuals, and are, therefore, more of a complementary set of findings to the age-based work.

The first two columns of Table 1 indicate the difference in economic losses – both by percent of GDP and dollar improvement relative to a base case – when we compare a uniform lockdown strategy, where at any one moment all age groups are locked down the same, with a targeted lockdown strategy, where different age groups can be locked down at different levels of severity. The results clearly show significant gains from a more targeted approach to lockdowns, with an 8.3 percentage point improvement in economic losses, which amounts to \$40 billion in savings (peak to trough).

Columns 3 and 4 show what happens when the targeted approach to lockdowns is taken and non-economic interventions are further improved. The improvement in economic losses over the targeted case is 5.7 and 5.2 percentage points, respectively, for a 5.0 percentage point improvement in testing and tracing and a 20 percent reduction in the rate of transmission, from, for example, better mask wearing. These gains relative to the uniform lockdown are \$67 billion and \$64 billion, respectively, or \$27 billion and \$25 billion above and beyond the basic targeted approach.

Figures 2 through 5 make clear how these savings come about. Figure 2 shows that, under a uniform approach, everyone must be locked down fully in the early stages (left-hand panel), leading to significant economic losses. The lockdown ends as the infection is stamped out (right-hand panel). In Figure 3, this is not the case, as only the older population must be locked down fully – and for even longer than under the uniform approach – with the younger and middle aged still locked down but with milder restrictions. Again, the lockdown ends abruptly for the younger and middle-age groups when the infection is stamped out. For the older group, the lockdown mostly ends as infections come under control, but not completely.

Figures 4 and 5, representing the improved testing and tracing and transmission, respectively, show a similar story as Figure 3, but exacerbated in that the older population

is locked down much the same, but the young and middle aged are locked down less severely, leading to additional economic gains.¹³

Overall, the far superior outcome in terms of economic losses, with no worsening of health outcomes, occurs when policymakers take a more targeted approach to economic interventions. Moreover, improvements in non-economic interventions result in further economic savings. We note, as described above, that nothing about these results would prevent a scenario where everyone is locked down the same for a short period, as is the case in parts of the country at the time of writing. Just because policymakers can choose to target does not mean they must; under the calibration we use, however, based on the data available, they do.

Table 1: Annual Economic Losses under Different Interventions

Economic Losses	Uniform Lockdown	Targeted Lockdown	Targeted with 5 percentage point Improved Testing and Tracing	Targeted with 20 percent Improved Transmission
Percent of GDP (%)	24.4	16.1	10.4	10.9
Improvement relative to base case (\$ billions)		40*	67	64

As an example, this calculation would be $((24.4-16.1)\text{Most recent GDP})/4$, with the denominator (4) turning an annual result into a quarterly one. We use a similar formula throughout our analysis. The most recent real GDP data at time of writing were from September 2020.

Source: Statistics Canada, table 36-10-0434-01.

¹³ The small rise in lockdowns at the end of the period in Figure 4 arises because infections are never removed fully. The models interpret this as an indication that herd immunity is not achieved, with the optimal response being an increase in lockdowns.

Figure 2: Uniform Economic Lockdown




Figure 3: Targeted Economic Lockdown




Figure 4: Targeted Economic Lockdown with Improved Testing and Tracing





Figure 5: Targeted Lockdown with Improved Transmission Rate



Industry-based Results

We see the industry results as a way to determine if a different approach to targeting would result in similar findings. In that sense, we think of these results as complementary.¹⁴

As Table 2 shows, targeting again produces better economic outcomes without sacrificing health outcomes, although the magnitude of these economic gains is lower than in the age-based results. Going from uniform to targeted lockdowns improves the economic loss by \$7 billion. Improving non-economic interventions further – including a 5 percentage point improvement in testing and tracing and a 20 percent reduction in the rate of transmission – leads to gains of \$26 billion and \$32 billion (peak to trough), respectively, over the uniform case, and to \$18 billion and \$25 billion improvements, again respectively, over the targeted approach.¹⁵

Overall, the gains are smaller because we are dealing with only the working-age population, leaving out the most vulnerable. Moreover, the higher-contact industries tend to be those populated by younger individuals, who are least at risk, so shutting them down does not significantly improve health outcomes. These higher-contact industries also suffer the greatest productivity loss from lockdowns.

Figures 6 through 9 give us clues as to what is happening, which is similar to the age-based results. In the targeted scenario, Figure 7, we see a lockdown of the low-contact industries at a lower rate than high-contact industries, unlike in Figure 6. Figures 8 and 9 show that, with mild improvements to testing and tracing and transmission, only the high-contact industries must be locked down.

Table 2: Annual Economic Losses under Different Interventions

Economic Losses	Uniform Lockdown	Targeted Lockdown	Targeted with 5 percent Improved Testing and Tracing	Targeted with 10 percent Improved Transmission
Percent of GDP (%)	11.4	9.9	6.1	4.7
\$ improvement relative to base case (<i>\$ billions</i>)		7	26	32

¹⁴ We also acknowledge that some high-contact industries include essential service businesses, such as groceries, cannot be shut down. Their inclusion within an industry category might affect the calibration and the exact economic savings from a more targeted approach, but we see no reason they would change the narrative that a targeted approach leads to smaller economic losses than a uniform approach (at no additional cost to health outcomes).

¹⁵ Rounding accounts for the discrepancy between \$18 billion and the \$19 billion one gets when subtracting \$7 billion from \$26 billion.

Figure 6: Uniform Economic Lockdown




Figure 7: Targeted Economic Lockdown




Figure 8: Targeted Economic Lockdown with Improved Testing and Tracing





Figure 9: Targeted Lockdown with Improved Transmission Rate



Variations to Non-economic Interventions

In both the age-based and industry-based results, we saw the gains from small improvements in non-economic interventions, specifically around testing and tracing and lowering transmission rates from, for example, better adherence to mask wearing. To reinforce the importance of these interventions, we now turn to showing how much better the economic loss would be if we made more significant gains – specifically, by improving our testing and tracing capabilities by 50 percent (or 30 percentage points, compared with just 5 percentage points above), and by cutting transmission by 50 percent (compared with just 20 percent above).

In the age-based case, the economic improvement in dollar terms over uniform lockdowns is \$89 billion for both the testing and tracing and transmission improvements (Table 3). This compares with \$67 and \$64 billion, respectively, for the more conservative improvements we used above. Figures 10 and 11 show that these improvements drastically shorten lockdown periods for all groups.¹⁶

Similarly, in the industry-based results (Figures 12 and 13), the improvement under each scenario is \$43 billion over the uniform lockdown case, compared with \$26 billion and \$32 billion, respectively, with the more conservative improvements above. Interestingly, there is essentially no lockdown of industries with this level of testing and tracing and transmission in place. We remind readers that the industry-results are only for working-age individuals and, therefore, this result should not be taken to mean there are no lockdowns anywhere in society. As we saw with the age-based results, which include older individuals, some degree of lockdown is still necessary.

Table 3: Annual Economic Losses under Different Targeted Non-economic Interventions

Economic Losses	Age-based 50 percent Improvement in Testing and Tracing	Age-based 50 percent Improvement in Transmission	Industry-based 50 percent Improvement in Testing and Tracing	Industry-based 50 percent Improvement in Transmission
Percent of GDP (%)	5.6	5.7	2.5	2.5
Improvement relative to base case* (<i>\$ billions</i>)	89	89	43	43

¹⁶ Readers might notice that \$89 billion is the entire peak to trough loss the Canadian economy experienced from February to April. This is because the baseline loss under the uniform strategy in our model simulations is \$116 billion peak to trough. It is greater than the actual peak to trough loss because the losses continue for longer under our simulation. As Table 3 shows, there remains a significant economic loss even under these much improved testing and tracing/improved transmission scenarios.

* The base case is relative to the age-based uniform lockdown (24.4 percent drop) and industry-based uniform lockdown (11.4 percent drop).

Figure 10: Age-based 50 percent Improvement in Testing and Tracing




Figure 11: Age-based 50 percent Improvement in Transmission Rate




Figure 12: Industry-based 50 percent Improvement in Testing and Tracing





Figure 13: Industry-based 50 percent Improvement in Transmission Rate



Much has been made about the different approaches the provinces have taken in fighting COVID-19. Although there are not enough available data to calibrate our entire model at the provincial level, we can show that not only are the absolute economic losses better when we have a better starting point in terms of capacity for testing and tracing, but the improvement in going from a uniform to a targeted approach is better in percentage terms as well.

We can use our baseline calibration above for both age-based and industry-based results, and assume this represents a large province such as Ontario, which, in the second wave, abandoned contact tracing in hot spots like Toronto. We can also take another large province, British Columbia, which did not, and set its testing and tracing success rate at, say, 15 percentage points higher than Ontario’s (a 25 percent improvement).

The first two columns of Table 4 just replicate the above results for the age-based modelling. The improvement when we go from uniform to targeting is approximately 34 percent $((24.4 - 16.1)/24.4)$. A better starting point in terms of testing and tracing improves first and foremost the economic losses under any lockdown approach. It also improves the gains in economic losses when going from a uniform to a targeted approach, with an approximately 43 percent improvement (columns 3 and 4). Comparing Figures 14 and 15 to Figures 2 and 3 above, the story is that these improvements shorten the most severe portions of a lockdown for all and, in the targeted case, while the older group is locked down more severely for longer, younger and middle-aged individuals return to normal more quickly.

Table 5 runs the same analysis for the industry-based results. We see the same gains in economic losses regardless of lockdown strategy when a province has a better testing and tracing starting point, but we do not see any gains when we go from a uniform to a targeted economic lockdown. As we see in Figures 16 and 17, however, which compare to Figures 6 and 7, the lockdowns are so mild in both the uniform and targeted cases that the gains to be had are just not that big – a good scenario to be in.

Table 4: Age-based Annual Economic Losses Based on Initial Starting Point, Ontario and British Columbia

Economic Losses	Ontario		British Columbia	
	Uniform Lockdown	Targeted Lockdown	Uniform Lockdown	Targeted Lockdown
Percent of GDP (%)	24.4	16.1	17.5	10.0
Improvement relative to base case (%)		34.0		42.6

Table 5: Industry-based Economic Losses Based on Initial Starting Point, Ontario and British Columbia

Economic Losses	Ontario		British Columbia	
	Uniform Lockdown	Targeted Lockdown	Uniform Lockdown	Targeted Lockdown
Percent of GDP (%)	11.4	9.9	3.2	3.2
Percent improvement relative to base case (%)		13.3		0

Figure 14: Age-based Uniform Lockdown, British Columbia




Figure 15: Age-based Targeted Lockdown, British Columbia




Figure 16: Industry-based Uniform Lockdown, British Columbia





Figure 17: Industry-based Targeted Lockdown, British Columbia



Robustness Checks

After running many robustness checks, we conclude, in general, that going from a uniform to a targeted approach to economic lockdowns lowers economic losses while not worsening health outcomes, with further gains to be had when non-economic interventions are improved.¹⁷

This conclusion is more robust in the age-based analysis than in the industry analysis, given that the former includes the entire population ages 20 and older, while the latter includes only working-age individuals – that is, it excludes the retired. As well, the most vulnerable individuals are working in the lowest-contact industries, meaning there are fewer health gains from focusing lockdowns on high-contact industries. These high-contact industries also face the greatest productivity loss when shut down.

One example where this divergence between age-based and industry-based results holds is when we ask what would happen if, instead of averaging the transmission rate across society, we assume a different rate for each of our groupings. We take this to be a supercluster/variant analysis of sorts, even if not a perfect one.¹⁸ Appendix C has more details on superclusters/variants, as well as the results from this robustness check.

¹⁷ For example, we looked at what happens with improved social distancing, where we lower our interactions with the most vulnerable. This and other robustness checks are available from the authors upon request.

¹⁸ See Tufekci (2020) for more on the importance of superclusters.

Conclusions

The usefulness of an epidemiological-economic model for policymakers is in simulating a wide range of possible interventions that achieve desired health outcomes while minimizing economic damage. The immediate relevance of these results is in providing policy insights as we move along the path toward widespread vaccine inoculation.

From our results, we offer the following insights relevant to the situation we are facing today:

- Targeted lockdowns in high-risk areas, such as the elderly or high-contact industries, are clearly preferable to uniform lockdowns. For example, the economic costs through targeting by age, rather than a uniform age lockdown, could reduce the economic losses – without worsening health outcomes – from an estimated \$116 billion to \$77 billion, a significant saving in terms of higher incomes, profits and taxes of approximately \$40 billion, or 8.3 percentage points of the peak-to-trough GDP loss.
- Non-economic interventions could further reduce the economic costs through improved testing and tracing or improved transmission rates from widespread mask wearing. When targeting by age, these non-economic interventions reduce costs by an additional \$27 billion and \$25 billion, respectively. The importance of such non-economic interventions cannot be overstated: they complement any lockdown strategy in further reducing economic costs; and they can sustain the progress achieved through lockdowns once they are removed, or even for the period before immunities are built up after vaccination.
- Critical to all the above are timely, accurate data, clear communications and investment in testing and tracing.

In addition to these general insights, the practical usefulness of our results can also be seen in applying them to specific situations in Canada today. For example, in the case of Ontario, which was forced to abandon testing and tracing in many hotspots, significant economic gains are still to be had from a targeted approach. A better starting point in British Columbia, however, which was able to continue to test and trace, markedly improved these gains. This explains why that province has been able to take an even more aggressively targeted approach than Ontario during the second wave.

We note again that nothing about the optimality of a targeted approach would prevent a complete lockdown of all non-essential activities, where everyone was locked down the same for a short period, as is the case in parts of the country at the time of writing.

In summary, a modelling exercise such as ours is but one tool in determining the appropriate policy response to COVID-19. We believe, however, that its ability to intersect health and the economy makes it of value to both health and economic policymakers as vaccines are rolled out over the coming year.

Appendix A: Deciding which Industries Fit in which Category

We use the industry breakdown from Statistics Canada's CANSIM table 36-10-0434-01 (GDP at basic prices, by industry, monthly). We keep it to the two-digit NAICS code, leaving us with 20 industries under consideration. As described in the text, we slot industries into the different categories using Statistics Canada data on the percentage of people within each industry who were able to work from home prior to the COVID-19 pandemic. Industries where the share was between 0 and 10 percent were considered high contact, those between 10 and 20 percent were considered medium contact and those above 20 percent were considered low contact. Many sub-sectors make up these different industry categories, so we were forced to average. The breakdown is as follows:

Agriculture, forestry, fishing and hunting [11] – HIGH CONTACT

Mining, quarrying, and oil and gas extraction [21] – MEDIUM CONTACT

Utilities [22] – MEDIUM CONTACT

Construction [23] – HIGH CONTACT

Manufacturing [31-33] – HIGH CONTACT

Wholesale trade [41] – MEDIUM CONTACT

Retail trade [44-45] – HIGH CONTACT

Transportation and warehousing [48-49] – HIGH CONTACT

Information and cultural industries [51] – LOW CONTACT

Finance and insurance [52] – LOW CONTACT

Real estate and rental and leasing [53] – LOW CONTACT

Professional, scientific and technical services [54] – LOW CONTACT

Management of companies and enterprises [55] – LOW CONTACT

Administrative and support, waste management and remediation services [56] – MEDIUM CONTACT

Educational services [61] – MEDIUM CONTACT

Health care and social assistance [62] – HIGH CONTACT

Arts, entertainment and recreation [71] – MEDIUM CONTACT

Accommodation and food services [72] – HIGH CONTACT

Other services (except public administration) [81] – HIGH CONTACT

Public administration [91] – MEDIUM CONTACT

Appendix B: Calibration Details

The list of parameters, a reminder of what they represent and the calibration broken down both by age group (young, middle aged and old) and industry (low contact, medium contact and high contact) are as follows:

- η – failure to detect and isolate an infected individual:
 - young, 0.6; middle aged, 0.6; old, 0.6 (source: Government of Canada and Canadian Blood Services and Canada’s COVID-19 Immunity Task Force)
 - low contact, 0.6; medium contact, 0.6; high contact, 0.6 (source: Government of Canada and Canadian Blood Services and Canada’s COVID-19 Immunity Task Force)
 - Acemoglu et al. (2020) assume at first an η of 0.9, meaning 90 percent of infected people are not detected and isolated. They then test a scenario where they reduce this failure rate to 0.7. Improvements in testing and tracing since their paper was written suggest we could use 0.7 at a minimum as our baseline here. We take it a step further and divide confirmed cases by seroprevalence (presence of antibodies), using the most recent data, which because of the denominator, forces us to use the end of June 2020. At this time, Canada’s confirmed case rate was about 0.275 percent (103,900 total cases), with the presence of antibody rate at 0.7 percent.¹⁹ Dividing these two numbers, we get about 40 percent detection, which we subtract from 100 percent to get 60 percent failure.
- γ – arrival of end of disease for non-ICU patients:
 - young, 1/18; middle aged, 1/18; old, 1/18 (source: Acemoglu et al. 2020)
 - low contact, 1/18; medium contact, 1/18; high contact, 1/18 (same assumptions as per age)
- β – equals $R_0\gamma$, where $R_0 = 2$ is the community transmission rate and $\gamma = 1/18$, as above:
 - young, 0.111; middle aged, 0.111; old, 0.111 (source: Acemoglu et al. 2020, with adjustments)
 - low contact, 0.111; medium contact, 0.111; high contact, 0.111 (same assumptions as per age)
 - Here we assume both by age and industry breakdown that the $R_0 = 2$. In Canada as of the end of August 2020, it was fluctuating tightly around 1, which

¹⁹ The seroprevalence rate is from the Canadian Blood Services and Canada’s COVID-19 Immunity Task Force.

is the level beyond which the infection rate begins to expand exponentially. It has obviously worsened since, and so we set $R_0 = 2$.

- θ – inefficiency measure for lockdown (for example, if 0.75, then lockdowns are 75 percent effective):
 - young, 0.75; middle aged, 0.75; old, 0.75 (source: Acemoglu et al. 2020)
 - low contact, 0.75; medium contact, 0.75; high contact, 0.75 (same assumptions as per age)
- ρ – interaction intra and inter-group:
 - young, 1; middle aged, 1; old, 1 (source: Acemoglu et al. 2020)
 - In the baseline, we assume every group interacts with itself and other groups the same and normalize to 1. We can then see the impact of tighter restrictions on interactions with the most vulnerable.
 - low contact, 0.76; medium contact, 0.94; high contact, 1 (source: Statistics Canada table 33-10-0228-01, looking at how much people worked from home by industry pre- and during COVID)
 - These are intra-group interaction rates. We also assume that inter-group interactions occur at the higher of the two groups. So, for example, the interaction level between a high-contact and a low-contact individual will equal the rate of the high-contact intra-group interaction.
 - Methodology: Calculate by industry the percentage of workers who worked from home pre-COVID. Take the average by industry type. Subtract this number from 1 to represent the contact rate for those forced to go into the “office.” Normalize to 1 the high-contact industries, which represents the highest level of contact possible. Adjust the medium- and low-contact industries accordingly.
- $\bar{\delta}$ – initial fatality rate:
 - young, 0.0013; middle aged, 0.0117; old, 0.2417 (source: Government of Canada, <https://health-infobase.canada.ca/covid-19/epidemiological-summary-covid-19-cases.html>, accessed September 2020²⁰)
 - Total deaths divided by total case count within the respective age bracket.

²⁰ Although these numbers obviously change every day, the direction has been largely the same since we accessed them, and is critical to the gains from a targeted approach. The data are broken down by deciles: 20–49 for the young, 50–59 for the middle aged and 60+ for the old.

- low contact, 0.0035; medium contact, 0.0022; high contact, 0.0020 (sources: Government of Canada: health-infobase.canada.ca/covid-19/epidemiological-summary-covid-19-cases; He, Messacar, and Ostrovsky 2017)
 - All three industry groups have a median age in the young age group set. However, all three medians are above the mid-point of the young group, so we assume the initial death rate by age group applies to the mid-point and make upward adjustments.
- $v_j - \sigma \bar{\delta}_j$, allows for modelling the effect of the infection rate on the death rate as a result of constrained ICU capacity (we assume that ICU needs are proportional to initial death rates), where $\sigma = 0.0076$ (source: Acemoglu et al. 2020)
- κ – the higher this number, the more infected people are detected and allowed to return to work when recovered:
 - young, 1; middle aged, 1; old, 1 (source: Acemoglu et al. 2020)
 - In other words, we assume we can identify all sick individuals and that they can return to work when recovered. We can also set it to 0 to test a complete inability to determine these individuals.
 - low contact, 1; medium contact, 1; high contact, 1 (same assumptions as per age)
- Δ_j – years (days) to retirement to measure the loss of economic activity from death:
 - young, $32.5 \cdot 365$; middle aged, $10 \cdot 365$; old, $2.5 \cdot 365$ (source: Acemoglu et al. 2020)
 - low contact, $21.4 \cdot 365$; medium contact, $25.5 \cdot 365$; high contact, $26.2 \cdot 365$ (source: He, Messan, and Ostrovsky 2017)
 - We take the median age by industry, subtract from an average retirement age of 67, and multiply by 365 to get total days to what would have been retirement.
- ω – earnings per capita:
 - young, 1; middle aged, 0.92; old, 0.3 (source: Statistics Canada, table 11-10-0239-01)
 - We normalize the young median income to 1, and adjust the others accordingly. We make an adjustment for the older group to account for those who do not earn wages.²¹

²¹ The available data only split earnings such that here the young are 16–54, the middle aged 55–64 and the old 65+.

- low contact, 0.86; medium contact, 1; high contact, 0.65 (source: Statistics Canada, table 14-10-0063-01)
 - We average weekly earnings by industry type. We then normalize the largest number to 1, and adjust the others accordingly.
- ζ – productivity rate at home:
 - young, 0.7; middle aged, 0.7; old, 0.7 (source: Acemoglu et al. 2020)
 - low contact, 0.76; medium contact, 0.69; high contact, 0.61 (source: Statistics Canada, table 33-10-0228-01, looking at how effective people are at producing today what they were producing pre-COVID-19²²)
 - We sum across different industry types the GDP left over despite restrictions keeping people at home and potentially shuttering businesses, then divide that by total GDP by industry pre-COVID-19.²³
- N – population percentage:
 - young, 0.53; middle aged, 0.26; old, 0.21 (source: Statistics Canada, table 17-10-0005-01)
 - The denominator in each case is total population for those ages 20+ (that is, the three groups)
 - low contact, 0.19; medium contact, 0.19; high contact, 0.62 (source: Statistics Canada, table 14-10-0291-01)

Appendix C: Supercluster/Variants Robustness Check

The ideal supercluster analysis would take on a stochastic instead of a deterministic nature to the transmission rate. In other words, in the first step, the initial transmission rate would be drawn at random, with some luck involved in the initial spread of the disease based on this draw. Think about countries where at first there were massive spreads versus

²² Note that these data would tell us how productive people were working from home in their respective industry, in most cases with children at home. As a check, given that children have returned to school, we could re-estimate with higher levels of productivity to represent the gains from not having children at home. However, this would not tell us anything about how many more children would get infected as a result of going to school, and perhaps increase infection rates in other groups. To handle this, we would have to increase contact rates between groups, assuming higher parents' infection rates from coming in contact with their children. We leave that for others to tackle in future research.

²³ One potential issue with this industry breakdown is that it has trouble capturing industries that did better during the lockdown period in March and April. However, only agriculture, forestry, fishing and hunting had a higher GDP at the end of April than at the end of February – and this increase was marginal. It also struggles to capture, over the course of the year, adjustments that made certain industries better off than they were pre-COVID as a result of industrial reorganization. This reinforces our use of peak-to-trough thinking regarding economic losses, as described above.

countries where there were not, yet where many of the features of the two countries, such as weather, demographics, economies and so on, are the same. And, instead of being the same throughout time, there would be other random draws across the simulation period. Our model cannot incorporate this stochastic nature, but we can allow for transmission rates to be higher in, say, the older age group or high-contact industries where there might be more exposure to poorly ventilated indoor spaces (for example, long-term care facilities, manufacturing warehouses). In some ways, we can treat this as a model of variants with higher degrees of transmission rates, which will disproportionately affect those that are at greater risk of exposure.

We assume a 20 percent greater transmission rate than in our primary specification in the case of the older group and high-contact industries. For the middle-age group and medium-contact industries, we use the same transmission rate as we used above. For younger individuals and low-contact industries, we assume a 20 percent lower transmission rate than in our primary specification.

The age-based results presented here tell a similar story in terms of improved economic losses, without worsening health outcomes, when we go from a uniform to a targeted economic lockdown (columns 1 and 2 of Table C.1). Whereas, with our initial results above, the improvement was from 24.4 percent to 16.1 percent of GDP, here it is from 25.1 percent to 11.3 percent of GDP. Indeed, the figures look similar as well (Figures C.1 and C.2), with a severe lockdown across all groups in the uniform case, before an abrupt end when infections are wiped out, and a severe lockdown for the oldest age group, but a lighter touch on those less vulnerable in the targeted approach. We note that the severe lockdown for the oldest group under the targeted approach is even longer than in the uniform case – similar to the baseline results.

The results are, indeed, similar, although the economic gain from going from a uniform to a targeted approach in this analysis is greater than in the baseline results. One explanation for this is that the health gains from locking down the elderly are much greater when the transmission rate is higher in this group – for example, because more of the elderly live in congregate settings. As a result, the older age group is locked down more severely for longer in the targeted case, and the younger groups – who contribute more to economic growth – are locked down less severely and for a shorter period.

In the industry case, we see no gain – indeed, a minor loss – from a more targeted approach. The lockdown design differs between the uniform and targeted approaches (Figures C.3 and C.4), but the economic outcome is quite similar. The lack of an economic gain when we go from a uniform to a targeted approach is because the higher transmission rate is in the high-contact industries, where the youngest workers are employed (so fewer health gains from lockdowns), and which face greater productivity loss from being shut down than workers in low contact industries.

Table C.1: Supercluster Robustness Check

Economic Losses	Age-based Uniform Lockdown	Age-based Targeted Lockdown	Industry-based Uniform Lockdown	Industry-based Targeted Lockdown
Percent of GDP (%)	25.1	11.3	12.0	12.4

Figure C.1: Age-based Uniform Lockdown, Supercluster




Figure C.2: Age-based Targeted Lockdown, Supercluster




Figure C.3: Industry-based Uniform Lockdown, Supercluster





Figure C.4: Industry-based Targeted Lockdown, Supercluster



References

- Acemoglu, Daron, Victor Chernozhukov, Ivan Werning, and Michael D. Winston. 2020. "Optimal Targeted Lockdowns in a Multi-Group SIR Model." NBER Working Paper 27102. Cambridge, MA: National Bureau of Economic Research.
- Ciuriak, Dan. 2020. "Who Knew What When: The International Transmission of Information on the COVID-19 Outbreak," Policy Brief, Asia Pacific Foundation of Canada, 25 August. Online at <https://www.asiapacific.ca/publication/who-knew-what-when-international-transmission-information>.
- He, Jiaosheng, Derek Messacar, and Yuri Ostrovsky. 2017. *The Relationship Between Firm Size and Age of Workforce: A Cross-Industry Analysis for Canada*. Statistics Canada Analytical Studies: Methods and References.
- Hedengren, John D., Reza Asgharzadeh Shishavan, Kody M. Powell, and Thomas F. Edgar. 2014. "Nonlinear Modeling, Estimation and Predictive Control in APMonitor." *Computers and Chemical Engineering* 70 (November): 133–48.
- Jenkins, Paul, and Jeremy Kronick. 2020a. "Dealing with the Second Wave." C.D. Howe Institute Intelligence Memo. October.
- . 2020b. "Dealing with the Second Wave: Further Results." C.D. Howe Institute Intelligence Memo. November.
- Kermack, William Ogilvy, A. G. McKendrick, and Gilbert Thomas Walker. 1927. "A Contribution to the Mathematical Theory of Epidemics." *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 115 (772): 700–21.
- Tufekci, Zeynep. 2020. "This Overlooked Variable Is the Key to the Pandemic." *Atlantic*, September 30. Online at <https://www.theatlantic.com/health/archive/2020/09/k-overlooked-variable-driving-pandemic/616548/>.
- Wächter, Andreas, and Lorenz T. Biegler. 2006. "On the Implementation of an Interior-Point Filter Line-Search Algorithm for Large-Scale Nonlinear Programming." *Mathematical Programming* 106: 25–57.