



CENTER FOR
**Brains
Minds+
Machines**

CBMM Memo No. 104

March 28, 2020

Can we Contain Covid-19 without Locking-down the Economy?

Shai Shalev-Shwartz and Amnon Shashua

Abstract

We present an analysis of a risk-based selective quarantine model where the population is divided into low and high-risk groups. The high-risk group is quarantined until the low-risk group achieves herd-immunity. We tackle the question of whether this model is safe, in the sense that the health system can contain the number of low-risk people that require severe ICU care (such as life support systems).



**This work was supported by the Center for Brains, Minds and
Machines (CBMM), funded by NSF STC award CCF-1231216.**

Can we Contain Covid-19 without Locking-down the Economy?

Shai Shalev-Shwartz and Amnon Shashua*

March 28, 2020

Abstract

We present an analysis of a risk-based selective quarantine model where the population is divided into low and high-risk groups. The high-risk group is quarantined until the low-risk group achieves herd-immunity. We tackle the question of whether this model is safe, in the sense that the health system can contain the number of low-risk people that require severe ICU care (such as life support systems).

1 Introduction

One could consider three models for handling the spread of Covid-19.

1. Risk-based selective Quarantine: Divide the population into two groups, low-risk and high-risk. Quarantine the high-risk and gradually release the low-risk population to achieve a managed herd immunity of that population. The managed phase is designed to allow the health system to cope with the expected number of severe cases. Given the herd immunity of the low-risk group, we can gradually release the high-risk population. The question is how to manage the release from quarantine of the low and high risk populations in a way that will not overwhelm the health system.
2. Containment-based selective quarantine: Find all the positive cases and put them in quarantine. This requires an estimation of $[t_0, t_1]$ the “contagious time interval” per age group, then given this time interval one could recursively isolate all the individuals at risk from a person that is carrying the virus using “contact tracing”. Another tool is predictive testing using contact-tracing to identify people with many contacts with other people and perform tests on them.
3. Country wide (or region wide) lock-down until the spread of the virus is under control. The lock-down could take anywhere from weeks to months. This is the safest route but does not prevent a “second wave” from occurring.

Models 2,3 could work in tandem and have been tried in China and Singapore. Model 3 is currently the default model around the globe and naturally has a tremendous crippling impact on the economy. In the remainder of this white-paper we derive some tools for analyzing the viability of the risk-based model. Specifically, what level of sampling and confidence level can be obtained to make sure that the health system can contain the model?

2 The Risk-based Selective Quarantine Model: How do we Know Whether it is Safe?

Consider a plausible definition of a high-risk group based on a cut-off age and certain pre-existing conditions. For the sake of concreteness, assume the cut-off age is 67+ which represents the retired segment of society. The low-risk group is the remainder of society which are released to their daily routine while following certain distancing protocols

*Both authors are affiliated with The Hebrew University of Jerusalem, Israel

that is aimed at slowing the spread, while keeping the economy un-disrupted to a large degree, but ultimately reaching a herd immunity level. At that point it is safe to gradually release the high-risk group from quarantine. The question is how do we guarantee that the health system will not be overwhelmed during the spread of the virus in the low-risk group?

Let b be the number of *severe* cases, e.g., those that require an intensive care unit (ICU), that the health system can handle, say $b = 600$ for a country of the size of Israel. Let m_q be the number of low-risk people that will develop severe symptoms and will require an ICU assuming we will adopt the risk-based selective model. Then, the model is “safe” if $b > m_q$. Our goal is to derive an upper bound on m_q , so that we’ll be able to ensure that $b > m_q$. Let m be the size of the low-risk group and let ν be the probability that a person that comes from the low-risk group will develop severe symptoms, assuming the person is currently sick. Then,

$$m_q = m \cdot \nu .$$

So, to ensure that $b > m_q$ we will require that $\nu < b/m$.

Before we continue, we note that fully understanding the dynamic of the spread of the virus and the dynamic of the development of the disease (for the sake of knowing when we will need ICUs and for how long) is very challenging and will probably require much more research and time. What we propose here is a *worst-case* analysis. The idea is to adopt a pessimistic view and show that even under this pessimistic view, the health system is not likely to collapse. We already made a pessimistic assumption since we did not take into account the facts that not all of the low-risk population will get sick, and even those that will get sick will not get sick at the same time and will not need an ICU at the same time.

Continuing our derivation, let p^* be the *current*, unknown, percentage of positive cases among the low-risk population and let k be the number of severe cases among the low-risk population from today until two weeks from now. Assuming that people that are positive cases will either develop severe symptoms within two weeks or will never develop severe symptoms, then we can estimate ν by $\frac{k}{p^* m}$. Note that again we are taking a worst-case view. Maybe some of the severe cases in two weeks will be due to people that are not infected today but will get infected tomorrow. This is even likely when the pandemic grows at an exponential rate. Nevertheless, to be on the safe side, we perform a worst-case analysis. Based on this notation, we can upper bound the value of ν as follows:

Lemma 1 Fix some $\delta \in (0, 1)$ and let

$$\tilde{k} = k + \sqrt{2 \log(1/\delta) k} + 4 \log(1/\delta) .$$

Then, with probability of at least $1 - \delta$ we have that

$$\nu \leq \frac{\tilde{k}}{p^* m} .$$

The proof of the lemma follows directly from measure concentration bounds (see for example Lemma B.10 in [2]).

Based on this lemma, we can upper bound m_q (with probability of at least $1 - \delta$) by

$$m_q \leq \frac{\tilde{k}}{p^* m} \cdot m = \frac{\tilde{k}}{p^*}$$

since even if *all* of the young population will get the virus, the number of severe cases will be the right hand side of the above equation. Of course, in reality we expect m_q to be much smaller than the above, both because not all of the young population will get the virus and because not all the severe people will be sick at the same time. Nevertheless, as mentioned previously we adopt a worst-case analysis.

The health system will be able to treat all of these severely sick people if $b > m_q$, from each we obtain the requirement:

$$p^* > \frac{\tilde{k}}{b} .$$

Since \tilde{k} and b are known, it is left to derive a method for ensuring that p^* is large enough.

This is a classical hypothesis testing problem: denote $\tilde{p} = \frac{\tilde{k}}{b}$, the null hypothesis is that $p^* \leq \tilde{p}$, and the alternative hypothesis is that $p^* > \tilde{p}$. By sampling n persons, uniformly at random from the low-risk population, we would like to find a cutoff value such that if the number of positive cases in our sample is above the cutoff, then we can reject the null hypothesis with a confidence of, say, 95%.

While one can find cutoff values for different values of n by using tables of the Binomial distribution, the following theorem specifies formulas for the sample size and for the cutoff value. One can verify by numerical calculations that the theorem is rather tight in the relevant regime.

Theorem 1 For $\tilde{p} \in (0, 0.1)$, setting the sample size to be $n = 4.438/\tilde{p}$ and the cutoff to be 10, we can reject the null hypothesis with a confidence of at least 95%. In other words, if we test $n = 4.438 b/\tilde{k}$ random people for the virus¹ and find out that there are more than 10 positive cases, then with probability of at least 95% we know that the risk-based model is safe.

The proof of the theorem follows directly from the following lemma.

Lemma 2 Fix $\delta \in (0, 1)$ and $\tilde{p} \in (0, 1/2)$. Suppose we sample $n = \frac{\frac{4}{3} \log(1/\delta)}{\tilde{p}(1-\tilde{p})}$ i.i.d. Bernoulli variables with parameter p^* and get at least $\left(\frac{4}{3(1-\tilde{p})} + 2\right) \log(1/\delta)$ positives. Then, with probability of at least $1 - \delta$ we have that $p^* \geq \tilde{p}$.

Proof If $p^* \geq \tilde{p}$ there is nothing to prove, so from now on we assume that $p^* < \tilde{p} < 0.5$. Denote by S_n the number of positive cases. Observe that S_n is the sum of n Bernoulli variables, each of which has probability of p^* to be positive. Then, Bernstein's inequality [1] tells us that for every $t > 0$,

$$\mathbb{P}[S_n - np^* > t] \leq e^{-\frac{t^2/2}{np^*(1-p^*)+t/3}}$$

Using $p^* < \tilde{p}$ we get that

$$\mathbb{P}[S_n - np^* > t] \leq e^{-\frac{t^2/2}{n\tilde{p}(1-\tilde{p})+t/3}}.$$

Setting

$$t = 2 \log(1/\delta) \quad , \quad n = \frac{\frac{4}{3} \log(1/\delta)}{\tilde{p}(1-\tilde{p})}$$

we have that

$$n\tilde{p}(1-\tilde{p}) = \frac{4}{3} \log(1/\delta).$$

Therefore,

$$\begin{aligned} \mathbb{P}[S_n - np^* > 2 \log(1/\delta)] &\leq e^{-\frac{t^2/2}{n\tilde{p}(1-\tilde{p})+t/3}} \\ &= e^{-\frac{(2 \log(1/\delta))^2/2}{\frac{4}{3} \log(1/\delta) + \frac{2}{3} \log(1/\delta)}} \\ &= \delta. \end{aligned}$$

Using again $p^* < \tilde{p}$ we also have that

$$\mathbb{P}[S_n > n\tilde{p} + 2 \log(1/\delta)] \leq \mathbb{P}[S_n > np^* + 2 \log(1/\delta)].$$

Combining all of the above yields

$$\mathbb{P}[S_n > n\tilde{p} + 2 \log(1/\delta)] \leq \delta.$$

¹Note that current tests are PCR tests and not serology tests, which only strengthen our results. Furthermore, PCR tests are known to have a negligible level of false positives.

This means that the probability that S_n will be at least

$$n\tilde{p} + 2 \log(1/\delta) = \left(\frac{4}{3(1-\tilde{p})} + 2 \right) \log(1/\delta)$$

is at most δ , which concludes our proof. ■

Next, we may wonder under what conditions would the hypothesis testing likely succeed? we can argue that if $p^* \geq 4.438\tilde{p}$ we are likely to see at least 10 positive cases. The reasoning is based on the following thread. The average positive cases we expect to see in a random population of size n is np^* , and the standard deviation is $\sqrt{np^*(1-p^*)}$. It follows that with high probability, the number of positive cases will be at least $np^* - 2\sqrt{np^*(1-p^*)}$ (we took two standard deviations). It is easy to verify that if $np^* \geq 19$ then $np^* - 2\sqrt{np^*(1-p^*)} \geq 10$. Hence, if $p^* \geq 4.438\tilde{p}$, we are likely to see at least 10 positive cases, and our validation of the risk-based model will succeed.

2.1 Estimating p^* from pooled tests

A pooled test is obtained by taking a sample from T persons, mixing all of these samples, and searching for traces of the virus. If the pooled test is positive, it means that at least 1 of the persons is positive. The following lemma generalizes Lemma 2 to pooled tests.

Lemma 3 Fix $\delta \in (0, 1)$ and denote $\tilde{p} = \frac{k}{b}$. Let T be a pool size and denote

$$\tilde{\phi} = 1 - (1 - \tilde{p})^T .$$

Suppose we sample $n = \frac{\frac{4}{3} \log(1/\delta)}{\tilde{\phi}(1-\tilde{\phi})}$ pools, uniformly at random from the low-risk population, test all of them for the virus, and find out that there are at least $\left(\frac{4}{3(1-\tilde{\phi})} + 2 \right) \log(1/\delta)$ positive pools. Then, with probability of at least $1 - \delta$ we have that $p^* \geq \tilde{p}$.

Proof Let ϕ^* be the probability that for a random pool of size T we will have $S_T \geq 1$, where S_T is the number of positive cases in the pool. Note that by definition of the pool test, it ends up positive if and only if $S_T \geq 1$. Using the same proof of Lemma 2, we have that for the n specified in the lemma, if we get at least $\left(\frac{4}{3(1-\tilde{\phi})} + 2 \right) \log(1/\delta)$ positive pools then with probability of at least $1 - \delta$ we have that $\phi^* \geq \tilde{\phi}$. Next, observe that

$$\phi^* = \mathbb{P}[S_T \geq 1] = 1 - \mathbb{P}[S_T = 0] = 1 - (1 - p^*)^T$$

Hence, if $\phi^* \geq \tilde{\phi}$ we have that

$$1 - (1 - p^*)^T = \phi^* \geq \tilde{\phi} = 1 - (1 - \tilde{p})^T .$$

Rearranging terms, we obtain that $p^* \geq \tilde{p}$. ■

The gain from the pool is clearly observed if $\tilde{p}T \ll 1$. In this case, using the approximation $1 - x \approx e^{-x}$ we have that $\tilde{\phi} \approx \tilde{p}T$. Hence, the number of pooled tests we need is roughly T times less than the number of regular tests we need, in order to show the same conclusions on p^* .

3 Discussion and Implications

In the event a risk-based quarantine approach would be contemplated by decision makers, the purpose of this document is to provide decision makers a formal and tight bounds to investigate whether the health system can cope with the number of severe cases that would reach ICU. Embedded in the reasoning is the idea of *selective* quarantine (based on age groups and existing pre-conditions, but could be any other criteria) where the "high-risk" group (the one we

suspect will have a high rate of severe cases) is quarantined and the other is allowed to spread the virus under certain distancing protocols. The underlying premise is that a full population wide quarantine is not a solution in itself — it is merely a step to buy time followed by a more managed (non brute-force) approach. The managed phase underlying our thinking is to create herd immunity of the low-risk group in a controlled manner while keeping the economy going. It is all about keeping the health system in check and not overwhelming its capacity to handle severe cases. The question we ask in this document is whether we can estimate in advance, through sampling, that the number of severe cases arising from the low-risk group would not overwhelm the system?

We conclude that this selective quarantine approach will work if p^* , the current percentage of positive cases among the low-risk (non-quarantined) population, is not very small, relative to $\tilde{p} = \tilde{k}/b$ which is roughly the ratio between the known number of severe cases from existing data and the capacity of the system (number of ICUs for example). Using concentration bounds we show that if we randomly sample $n \approx 1/\tilde{p}$ people from the low-risk population and find at least few positive cases, then the selective quarantine approach will succeed. Lemma 2 provides the precise bound with a $1 - \delta$ confidence. In Lemma 3 we extend the result to include “pooled” sampling as well. We also conclude that if we have other means of estimating p^* we can use the requirement $p^* \geq \tilde{k}/b$ in order to estimate the needed capacity b . For example, random sampling of a population of 2000 individuals in Iceland indicated that $p^* \approx 0.01$ from which we can deduce that $b \approx 100\tilde{k}$ ICUs should suffice for the low-risk population.

A critical element (besides the prevalence p^*) in our model is the random variable k which represents the number of severe cases among the low-risk group. The amount of medical resources b monotonically increases with k , thus, for an effective cut-off between low and high risk it would be very useful to obtain information (which today is missing) about the people that require intensive life support such as age distribution and distribution over pre-existing conditions in order to fine-tune the definition of the low-risk group such that k would be as small as possible.

It is worth noting that we focused on what is “safe” for the health system in the sense of how to estimate the number of severe cases that would be low enough not to overwhelm the ICUs. We ignored the fact that some severe cases could end up in the mortality statistics even when given proper care. In fact there are two probabilities to estimate (i) the probability of being in the “severe” category among the low-risk group, and (ii) the probability of mortality given proper care. We have bounded the former and ignored the latter. The reason for doing so is that the latter is beyond the scope of this paper because it is essentially a moral tradeoff between “safety” and “usefulness” that is employed in every aspect of society. For example, society does not put a lockdown on passenger car use in order to significantly reduce car accidents even though such a lockdown will save lives. Likewise, governments do not allocate infinite budgets for the health system even though there is a correlation between increased investments and saving lives.

As a final remark, going out of quarantine is a choice not an obligation. This is no different than people that are afraid of flights and decide not to go on an airplane. Families can decide to stay quarantined either as an extra safety measure or if some members of the family are from the high risk group while the others are from the low risk group.

Acknowledgements

The authors wish to thank the faculty members of the computer science department of the Hebrew University, as well as to Prof. Peter Bartlett, Prof. Nir Friedman, Prof. Katrina Ligett, Prof. Nati Srebro, Prof. Herve Bercovier, and Dr. Renana Eitan for comments and feedback on earlier drafts of this paper.

References

- [1] S. N. Bernstein. On a modification of chebyshev’s inequality and of the error formula of laplace. *Annals Science Institute Sav. Ukraine*, Sect. Math. 1, 1924.
- [2] S. Shalev-Shwartz and S. Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.