Optimal Social Distancing and the Economics of Uncertain Vaccine Arrival

Terrence Iverson, Larry Karp, and Alessandro Peri

The current, increased optimism about the imminent arrival and distribution of a COVID vaccine in the U.S. leads to an increased optimal level of social distancing. Under current circumstances, a policy that seeks to protect the most vulnerable, while allowing the bulk of the population to carry on without restrictions would be disastrous; that policy might have been loosely justified early in the pandemic. An early vaccine arrival time, in general, has an ambiguous effect on optimal social distancing. However, we find that currently in the U.S., the early arrival time promotes stricter social distancing.

Amid the surging COVID-19 pandemic, governments over the past year have faced a difficult policy choice: How much, and for how long, should they restrict economic activities that allow the virus to spread? A stricter social distancing policy—stronger quarantine requirements or the closure of more public spaces—slows the spread of the disease but inflicts economic costs on society. Optimal social-distancing policies should balance the economic costs of restricting mobility against the resulting benefit of reduced mortality costs.

The many sources of uncertainty about COVID and the consequences of social distancing complicate the government's policy problem. A crucial source of uncertainty concerns the "arrival time" of a vaccine, defined as the time needed to develop and distribute a vaccine widely. With this definition, vaccine arrival ends the need for social distancing. Expert opinion concerning the arrival time of a vaccine changed rapidly during 2020. Early in the pandemic, most experts viewed a one-and-a-half to three-year horizon as plausible, and some worried that COVID-19 could turn out to be a disease, like AIDS or the common cold, for which a vaccine would prove elusive. In the end, the U.S. approved a vaccine just eleven months after the start of the pandemic, though there remains uncertainty about the amount of time needed for widespread distribution.

We show how optimal social distancing policy depends on beliefs about the vaccine's arrival time. We ask how policymakers should respond to increased optimism about the early arrival of a vaccine. Should this optimism make social distancing policy stricter or less strict?

Mathematical models can discipline our attempt to answer such questions, helping to distinguish internally consistent and empirically grounded analysis from a merely plausible story. Fortunately, we have as our starting point a widely accepted model of disease contagion from epidemiology: the Susceptible-Infected-Recovered (SIR) model. This model, which is now routinely used by many economists working on COVID policy research, describes how individuals move from being initially susceptible to the virus, to becoming infected, to recovering—possibly, but not necessarily with immunity—or dying. We modify this familiar framework by treating the vaccine's arrival time as a random variable.

For COVID in the late 2020 and early 2021 U.S. context, we conclude that greater optimism leads to an optimal policy involving stronger social distancing. More generally, we show that greater optimism tends to increase optimal social distancing under a variety of circumstances (e.g., high costs of social distancing) that lead to high levels of infection. However, in circumstances that lead to low levels of infection, greater optimism tends to reduce the optimal level of social distancing.

We also evaluate proposals to move quickly to herd immunity. We focus on a best-case version of the Great Barrington Declaration (https://gbdeclaration.org/), which attracted considerable interest from the Trump administration in late fall 2020. The recommendation has two key features. First, it involves "targeted" instead of "uniform" policy. The former requires more vulnerable groups to adhere to stricter social distancing; uniform policy, in contrast, imposes the same level of social distancing on all groups, regardless of their risk factors. Second, it proposes that the vast majority of the population, those at lower risk, not be subject to any form of social distancing. The first feature is uncontroversial in principle, although there have been questions about its practicality. The second feature has attracted considerable controversy. Should society really require no social distancing for the bulk of the population?

When vaccine arrival remains far off, this proposal can be defended. Indeed, if mean vaccine arrival time is two years, the proposal performs almost as well as optimal uniform policy; however, it still leads to increased economic plus mortality costs of about $2 trillion compared to optimal targeted policy. Thus, early in the pandemic, when there was tremendous uncertainty and significant pessimism about the vaccine's arrival time, the proposal was plausible. But, if vaccine "arrival" is imminent—as in late 2020 and early 2021—then the policy is catastrophic. For the U.S., with a mean arrival time

Figure 1. Relationship Between Average Infection Rates and Expected Arrival Time

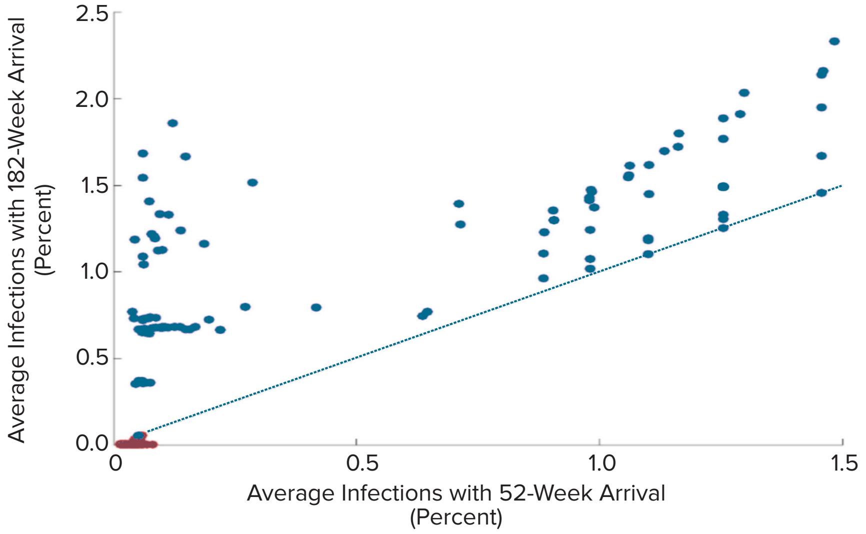

* Notes: Dots indicate average infection rate over the first year under two assumptions about mean arrival time: 52 weeks on the horizontal axis and 182 weeks on the vertical axis. The blue color indicates that a later expected vaccine arrival lowers optimal social distancing. The red color indicates the opposite.

of six months, 520,000 more of the vulnerable group (65 and over) die, 380,000 more of the less vulnerable group (under 65) die, and aggregate (economic + mortality) costs exceed $3 trillion. The assertion in the Great Barrington Declaration, that moving quickly to herd immunity will protect the vulnerable, is disastrously wrong: the vulnerable cannot be protected when infection levels are allowed to run extremely high.

U.S. policymakers care primarily about directly relevant policy advice, but economists are also interested in the underlying logic. Understanding this logic requires that we understand why, in general, the relation between social distancing and the vaccine's expected arrival time is ambiguous. The explanation turns on the fact that social distancing has two opposing effects. Stricter social distancing in the current period lowers the number of people who become infected in this period, thereby lowering the stock of infected people but also raising the stock of susceptible people in the next period. The reduction in the stock of infected people is a benefit to society because it lowers the probability of future infections, but the increase in the stock of susceptible people is a liability because those people remain vulnerable to future infection.

We refer to the future benefits associated with a reduction in the stock of infected people as the "infection channel" and the future liability associated with an increase in susceptible people as the "susceptible channel." The infection channel causes an earlier expected arrival time to lower optimal social distancing, while the susceptible channel has the opposite effect. The intuition for this result is that an earlier expected arrival time reduces the probability that society actually incurs the liability associated with the higher stock of susceptible: There is less time for those people to become sick with an earlier vaccine arrival. This reduction makes stricter social distancing relatively more attractive. Thus, when the susceptible channel is strong, an earlier expected arrival time increases optimal social distancing.

The net policy effect of the vaccine's uncertain arrival depends on the relative strength of these two channels; the relative strength depends on parameter values that reflect the actual or estimated U.S. experience and on current stocks of susceptible and infected populations. Important parameters include those that determine the economic cost of social distancing, the efficacy of social distancing, reinfection rates, and the transmissibility of the disease. For example, a high economic cost from reduced labor supply or a low valuation of mortality both lead to an optimal path for the economy in which infections are rather high. With high levels of infection, the susceptible channel dominates: here, the high stock of infected raises the probability that a susceptible person becomes infected, thereby raising the cost of having one more person being vulnerable to infection. In this setting, an earlier expected arrival time increases optimal social distancing because society is less worried about the potential liability being realized.

Our best estimates of parameter values and current stock levels place the United States in late 2020 and early 2021 into this category, where the susceptible channel dominates. As noted above, this conclusion implies that U.S. policymakers should restrict economic activities now more than ever because widespread distribution of a vaccine is imminent. Proponents of a rapid move to herd immunity (e.g., signatories of the Great Barrington Declaration) want to follow the opposite path.

But there are also plausible settings, e.g., with relatively low social distancing cost or relatively high mortality cost, where an earlier expected vaccine arrival time makes optimal social distancing policy weaker. In these circumstances, the optimal path of the economy keeps infections low, and the infection channel dominates. In the U.S., parameters consistent with this relationship may have been plausible early in the pandemic, especially if policy response had been effectively coordinated at the federal level. Now that U.S. infections are widespread and to a large degree out of control, that situation is implausible. In countries

Figure 2. Effect of a 10% Increase in Social Distancing on Mortality (left) and Stocks of Susceptible and Infected Populations (right)

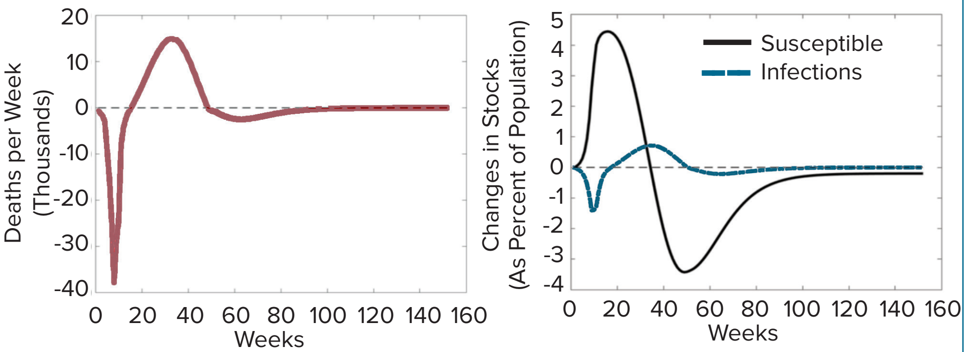

Impulse response functions associated with a 10% increase in initial-period social distancing relative to the optimal path.

* Note: Simulations for our baseline model. The mean arrival time is 156 weeks.

where infection levels have been consistently kept low, including Vietnam and New Zealand, this description of the world is realistic.

Our paper uses both analytical and numerical models to develop the intuition above. Figure 1 illustrates a key feature of the relationship between the level of social distancing policy and the expected time of vaccine arrival. We numerically solve our quantitative model for over 1,000 combinations of parameters and for two expected vaccine arrival times, 52 weeks and 182 weeks. Each dot in Figure 1 shows, for a single combination of parameters, the first-year average infection level when the expected arrival time is 52 weeks (horizontal axis) and 182 weeks (vertical axis).

All of the blue dots lie above the 45-degree line, indicating that for the corresponding parameters, a later expected arrival time lowers optimal social distancing, thus increasing infections. The red dots lie below the line, indicating that a later expected arrival time increases optimal social distancing, thus lowering infection. The red dots are all very close to the origin; they correspond to parameters that maintain low infection levels. These results show that an earlier expected arrival time increases optimal social distancing in circumstances that lead to high levels of infection (the blue dots).

The left and right panels of Figure 2 trace the impact of a 10% increase in early social distancing relative to the original optimum, holding future policy fixed. Economists call these graphs "impulse response functions." This figure corresponds to the case where the susceptible channel is strong so that earlier expected vaccine arrival makes optimal policy stricter. The left panel shows that the stricter social distancing substantially lowers deaths in the short-run, followed by a smaller but longer-lasting increase in deaths.

To explain the dynamics, the right panel shows that while infections initially fall (the blue curve), the stock of susceptible rises (the black curve). Individuals who avoid infection in the short run also avoid the benefit of developing immunity and hence remain in the susceptible pool. A higher stock of susceptible is fuel for the fire of future infections. Some of this fuel ignites in subsequent periods, leading to a resurgence in infections and deaths.

The uncertainty early in the pandemic about when a COVID vaccine would be developed and distributed made it hard to determine the optimal level of social distancing. Throughout the pandemic, there were calls to move rapidly to achieve herd immunity; some claimed that such a policy would result in fewer deaths among the most vulnerable populations over the long run. Most public health experts were extremely critical of those proposals, although they attracted considerable interest by the Trump administration. By late 2020, the prospects for a vaccine were much brighter, and in December 2020, the U.S. approved the first one. In early 2021, there remains some uncertainty about the timing of widespread vaccine distribution, but this uncertainty is hugely diminished.

We have experienced increased optimism over time about the imminent "arrival" of a vaccine. But it is easy to imagine a world in which the scientific problem of developing the vaccine proved harder than originally thought. In that case, we would have become less optimistic over time. In general, the relation between the optimal level of social distancing and optimism about vaccine arrival depends on many features. For the current U.S. context, our research strongly suggests that increased optimism should be met with stricter social distancing.

Suggested Citation:

Iverson, T., L. Karp, and A. Peri. 2020. "Optimal Social Distancing and the Economics of Uncertain Vaccine Arrival." *ARE Update* 24(3): 9–11. University of California Giannini Foundation of Agricultural Economics.

Authors' Bios

Terrence Iverson is an associate professor of economics at Colorado State University, Fort Collins. Larry Karp is a professor in the ARE department at UC Berkeley, who can be contacted at karp@berkeley.edu. Alessandro Peri is an assistant professor at University of Colorado at Boulder.

For additional information, the authors recommend:

"Optimal Social Distancing and the Economics of Uncertain Vaccine Arrival." Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3754054.