Comments and Discussion

COMMENT BY

DARON ACEMOGLU It is painful in the extreme to see that the extraordinary talents of Emmanuel Farhi, which are amply reflected in this paper, have since been lost to our profession. Emmanuel Farhi was a brilliant mind and a kind and generous person who illuminated, inspired, and influenced all of us and still had so much to contribute to our profession and the policy world. It is with a heavy heart that I am completing this comment on what is probably one of his last papers. He will be missed by our profession and of course by his many loving friends and collaborators among us.

As this comment is going to press, COVID-19 cases in the United States have already exceeded 25 million, deaths have surpassed 421,000, and the second wave that started in the fall has turned out to be much worse than most people expected. Any guidance for policy (if anybody were to listen) would be welcome in these troubling times. The paper by Baqaee, Farhi, Mina, and Stock is a timely and important contribution, providing exactly this type of guidance.

The authors develop an extended SEIR (susceptible, exposed, infected, and recovered) model that features five age groups arrayed in sixty-six sectors and contact matrices that depend on age and sector. The model is expertly calibrated, and the authors use it to evaluate various reopening scenarios.

The bottom line of the paper is simple and powerful: smart reopening policies can save both lives and the economy. The authors also give guidelines on what smart reopening should be. The most important lesson from this careful analysis is that economic measures, such as lockdowns, are neither necessary nor sufficient to control the pandemic. Instead, noneconomic nonpharmacological interventions, including personal distancing, face masks, and limits on large gatherings, can be very effective while

allowing the economy to return to some normalcy. The authors additionally (in my opinion very rightly) stress the importance of protecting the most vulnerable, in particular the elderly.

There is no doubt in my mind that if policymakers listened to and engaged with the authors, the United States would benefit significantly (though I am not optimistic about policymakers doing so, unfortunately). My very positive assessment notwithstanding, it is useful to place this paper in a broader context and point out areas for improvement, mostly for future research by the authors as well as other scholars.

BROADER CONTEXT How does the paper relate to the prior body of work? Two literatures should be distinguished in answering this question. The first is the voluminous literature in epidemiology using SIR (susceptible, infectious, and recovered) models, and especially the subbranch that focuses on COVID-19. A well-known example is the early work by Ferguson and others (2020). These works are very detailed in terms of heterogeneities in the population and infection dynamics, though less satisfactory when it comes to cost of lockdowns (because they do not model economic interactions). The authors of the current paper break new ground relative to these works by introducing a richer modeling of the economic aspects—for example, by recognizing differences across sectors in social contacts and costs of lockdown.

The second literature starts with more parsimonious models of infection dynamics, like the original SIR model of Kermack and McKendrick (1927), and aims to derive general lessons about how epidemics come to an end, how damaging they are, and how to deal with them. Early COVID-19 works in economics, starting with Atkeson (2020) and Alvarez, Argente, and Lippi (2020), have contributed to this literature. My own work in this area (Acemoglu and others forthcoming), which is related to this paper by Baqaee and his colleagues, is also in this genre. It can be viewed as a special case of the current paper, since it focuses on an SIR model with three age groups and no sectoral structure (rather than five age groups and sixty-six sectors). There is only one sense in which our work is not a special case, which I discuss next.

THE ISSUE OF OPTIMAL POLICY Our paper (Acemoglu and others forthcoming) considers optimal policy, while Baqaee and his colleagues eschew this focus and introduce a type of feedback rule that may be followed by governors as a function of COVID-19 cases and unemployment in their states (though the authors also note that these feedback rules may stand for other constraints on policy). They then discuss general strategies for reopening (including noneconomic and nonpharmacological interventions

as well as guidelines on the speed of economic reopening) that have to be evaluated taking the policy feedbacks from the governors as given.

I believe that focusing on optimal policy helps us clarify the trade-offs. This is a particularly relevant issue for me, since I have spent much of my career studying problems of political economy, which are instances where actual policy choices have little to do with optimal policies. And unless we understand the political incentives and constraints facing collective choices, we would reach misleading conclusions and develop a faulty understanding of the situation.

I have not faltered in my commitment to political economy, but in this case there is an important benefit from characterizing optimal policy. This is for three reasons. First, the problem of controlling an infection is a complex one, and thus focusing on constrained solutions without understanding what would work best could lead to incomplete or even incorrect inferences. The possibility of a second wave illustrates this issue clearly. Ferguson and others (2020) saliently mentioned the likelihood of a second wave. But what is not clear from their analysis is whether this is inevitable because anything that would prevent a second wave would be more costly or whether a second wave would represent a policy failure. Without understanding the answer to this question, it is difficult to know how to approach and prevent a second wave.

Second, as our analysis in Acemoglu and others (forthcoming) highlights, the form of optimal policies differs greatly depending on the preferences of policymakers in terms of economic losses versus lives lost. In such situations, providing a transparent menu of choices is often most informative, and optimal policy analysis achieves this aim. This is what our work does by making use of "Pareto" frontiers between the two major objectives: saving lives and saving the economy.

Third, our paper also establishes that the form of optimal policy is simple, hence ameliorating concerns that it may be too complex to implement. In particular, we show that significant gains—of the order of one-third reduction in economic costs for the same cost in terms of lives lost or vice versa—can be achieved by simply having a differential lockdown on the more vulnerable, 65 and older age group versus the rest of the population.

Can Baqaee and his colleagues' choice of focusing on general reopening strategies while imposing a specific policy feedback by governors be justified? One way to do that would be to argue that governors will indeed follow such a feedback rule and that optimal policy that cannot be "implemented" via incentive-compatible choices of governors is irrelevant.

I believe this argument has merit, but I am not entirely convinced by it. If the authors' paper has a powerful message on the form of optimal policy, this could change governors' responses, especially if optimal policy is proven to be simple and significantly superior to those that would follow from the feedback rules of governors.

A second justification for Baqaee and his colleagues' focus on constrained, suboptimal policies may be that, as they rightly note, there is a tremendous amount of parameter uncertainty, thus imposing some specific values on transmission rates, mortality rates, contact rates, and individual social distancing; then computing optimal policies may not be very informative. This point is valid and important. However, the same parameter uncertainty also makes nonoptimal policy analysis similarly fragile, and one way to deal with it is to systematically investigate the robustness of optimal policies (or other analysis) to changes in parameters or parameter instability. There are other options, as well, that can be fruitfully pursued in future work. For example, instead of deterministic optimal control, one could investigate robust control, which would take parameter uncertainty and instability explicitly into account (Zhou and Doyle 1998; Hansen and Sargent 2008).

CONCEPTUAL LESSONS Another issue worth discussing is whether the framework the authors build generates new qualitative perspectives. To be sure, this is a high bar, partly because the authors' main aim is to carry out a detailed quantitative analysis (and they do this quite well). Nevertheless, I believe it is fair to ask this question because they are part of an emerging literature marrying economics and epidemiology, and the main fruits of this endeavor should not just be quantitative estimates in some specific instances, but also new conceptual lessons.

Fortunately, there are two ways in which the authors provide such lessons. I will next point out how more can be done in each case.

First, in investigating which sectors should be open first, the authors propose a powerful measure to focus on: GDP-to-risk ratio, defined as how any policy variation (e.g., opening one of the sectors) will have an impact on GDP relative to its effect on the basic reproduction rate of the virus (the so-called R_0), which measures how many new infections one more infected individual will generate. The GDP-to-risk ratio is a useful measure, but it could be further developed. First, additional illustrations of whether ranking sectors according to GDP-to-risk ratio does better than alternatives would have been useful. Second, and more fundamentally, a more systematic analysis of what this measure captures would also be useful to undertake in future work.

In essence, the GDP-to-risk ratio is not a sufficient statistic. One way of seeing this is as follows. The basic reproduction rate, R_0, is the largest eigenvalue of the (linearized) dynamical system defined by the model.[1] The largest eigenvalue is informative about the speed of convergence of the dynamical system induced by the multigroup SIR model considered by the authors (e.g., Draief and Massoulié 2009, theorem 8.2). Yet it does not provide sufficient information about the mixing properties of the dynamical system. These mixing properties, which loosely speaking capture how quickly the infection jumps from one part of society to another, are critical for the spread of the virus. To see this we can consider a simple example. Suppose the multigroup setting approximates what is called an "island model" in the context of social and economic networks (Jackson 2010), whereby an individual on an island has a very high probability of interacting with others from the same island, and a small probability of interaction with those from other islands. Suppose, for the sake of this example, that each one of the n groups has approximately the same overall contact rate, but some groups have higher rates of contact with other islands. In this setup, increasing infection in any one of the n groups will lead to the same change in R_0 (or in the largest eigenvalue). Nevertheless, confirming that R_0 does not contain all the relevant information, the consequences of infection in a group that has higher interactions with the rest are worse. The second largest and other eigenvalues are informative about this type of mixing between different groups or islands. This example illustrates how just relying on R_0 or the largest eigenvalue is not sufficient for understanding the dynamics of the infection. Along these lines, one could either develop a more comprehensive measure than the GDP-to-risk ratio or investigate in the context of the quantitative exercise how much we are missing by focusing on the GDP-to-risk ratio.

A second way in which the analysis can be pushed further is by evaluating whether all of the heterogeneity is critical, or if one could have reached the correct conclusions with a stripped-down model. This would be particularly important for future work, which will likely build on the current paper's insights regarding which simplifications are justified when other related aspects of the questions are being explored. For example, what happens if the authors used three age groups and a single sector as in Acemoglu and others (forthcoming) instead of five age groups and sixty-six sectors? In what ways would this simplification lead to quantitatively or qualitatively

1. More precisely, it is the largest eigenvalue of the next-generation Jacobian matrices. See Diekmann, Heesterbeek, and Metz (1990).

different conclusions? A discussion of these issues would have helped future work in the area by conceptually clarifying which dimensions of heterogeneity matter more and for what reason.

Let me again give a brief discussion of this issue by drawing on Acemoglu and others (forthcoming). In that paper, the main (qualitative) conclusion is that "semi-targeted" policy that involves a different lockdown on the 65 and older age group can significantly improve either economic outcomes or public health outcomes or both. For example, in our baseline parameterization, a move from uniform policies to semi-targeted policies can reduce economic damages by one-third without any more lives lost. Additional targeting within the younger groups has very few benefits. This has a good economic reason. The 65 and older age group is the most vulnerable to the virus, with a case fatality rate of about sixty times those between the ages of 20 and 50. To save lives, this group needs to be protected from the infection, and because different groups interact frequently and even strict lockdowns are not perfect, just locking down this older group is not sufficient. Semi-targeted optimal policy, instead, also imposes a relatively lengthy lockdown on other age groups in order to reduce their infections and the transmission of infection to the older group. This explains why finer targeting is not very effective: it can trade off infections among those in their fifties against infections of those younger than 50, but this has little mortality benefit, and both groups will transmit the virus to those older than 65. Hence, it is approximately optimal to use a semi-targeted policy that applies a strict lockdown (a "protective custody" so to speak) on the 65 and older group and, simultaneously, moderates the spread of the infection among the rest of the population. A slower growth of infections among the younger groups is critical to ensure that the virus does not spread to those over 65 (e.g., because of interactions within households, in nursing homes, or in the context of other inevitable interactions).

It would be interesting to investigate more systematically to what extent these conclusions remain true when there are more age groups and also sectoral differences. This is particularly important, in my opinion, because new work in the emerging epidemiology-economics intersection is most promising if it simultaneously provides general qualitative lessons and realistic quantitative evaluations. The current paper is an excellent example of the latter, but attempting to draw out some of the general lessons would enrich it and improve the evaluation of the quantitative contributions.

CHOICES The authors should be commended for their state-of-the-art careful parameterization that captures the most salient aspects of the

COVID-19 outbreak and its economic costs. It is for this reason that it has the potential to have an impact on the policy debate.

One place where the paper, and I hope future papers in this area, can improve their exposition is in justifying and motivating their authors' choices. For example, the model is very detailed in terms of its sectoral structure and fairly granular when it comes to age groups. It also engages with the education sector, which is of course a key for a broad reopening of the economy, since parents of school-age children cannot be fully reintegrated into the labor market when their children stay at home.

Four choices of the authors deserve more discussion. First, while the model incorporates 66 sectors, it does not model the input-output linkages between sectors and approximates the contribution of each sector by a Hulten-like first-order approximation (Hulten 1978), ignoring any complementarities between sectors that may become important when some open up while others remain largely closed. These choices may or may not be important, and it would be useful for the authors to explore these issues (or at the very least do more to justify these choices). For instance, a sector that has very low contact may appear as a good candidate for early reopening, but if it has critical inputs from a high-contact sector, this might change the relevant calculus. Or keeping certain sectors, such as retail, closed might significantly reduce the marginal contribution of other sectors, such as wholesale. The framework here is already detailed and flexible enough to explore these issues.

Second, and perhaps more important from a policy angle, the paper does not model infections within nursing homes, where about 40 percent of the US COVID-19 deaths have so far occurred (Chidambaram, Garfield, and Neuman 2020). This may be because doing so may have required some important extensions of the framework. Nursing homes are not only places where residents are highly vulnerable, but they also lead to the fast spread of the virus. This may require the introduction of the more heterogeneous network structure. Though this is clearly beyond the scope of the current paper, I would like to flag it as an important area for future research.

Third, the paper takes the contact matrix from the POLYMOD data set (Mossong and others 2008). One issue with these data is that they suggest fairly low contact rates for older individuals. This has major implications. For example, if we take these contact rates to be much lower than for other groups, then even without lockdowns the more vulnerable, older individuals would become infected at a much lower rate. This would then

reduce the need for strict lockdowns on this group. However, low infection rates for the older group appears counterfactual (certainly given the very high infection rates in nursing homes, which as I have already mentioned seem to suggest greater, not lower, infection rates for the subpopulation). There are good reasons for conjecturing whether POLYMOD may be missing important context for the elderly within families or in other contexts. Indeed, the more recent (and highly systematically collected) BBC Pandemic Data, which we used in Acemoglu and others (forthcoming), has significantly higher contact rates for the elderly. Given the importance of this issue, more discussion and more robustness checks would be useful.

Finally, as I mentioned already, I agree with the authors' emphasis on the importance of noneconomic, nonpharmacological policies. Face masks appear to be critical for reducing infection on the basis of existing evidence and research on the trajectories of droplets (Chu and others 2020; Greenhalgh and others 2020). The paper and the authors should be commended on emphasizing and highlighting the major economic and public health benefits obtainable from noneconomic, nonpharmacological policies. More variations and robustness checks on this issue, especially a more detailed discussion of whether face masks can reduce transmissions in various different sectors of the economy, could have been useful.

OTHER ISSUES I would like to end this discussion with a brief mention of several issues that clearly lie beyond the scope of the current paper, but may be worth speculating on briefly.

There is every possibility that this pandemic may turn out to be what James Robinson and I called a critical juncture in our book *Why Nations Fail* (Acemoglu and Robinson 2012): an episode where existing institutions and social arrangements prove to be inadequate and thus it paves the way for major changes. We argued that, during such periods, small details matter and the direction of change is generally difficult to ascertain. If so, as important as dealing with the fallout from the current pandemic will be to prepare for what types of economic, political, and social changes will come. Though clearly not the focus of the current paper, pandemic-fueled social change is something the economics profession should start thinking about and may have useful ideas to contribute to.

Second, the economic costs of the pandemic have been lessened by digital technologies and automation that have enabled many sectors to function during lockdown. But it is also likely that the pandemic will give an additional boost to digital technologies and platforms as well as to efforts to further automate the economy. My recent work with

Pascual Restrepo (Acemoglu and Restrepo 2020a, 2020b) highlights the possible costs of investing too much in automation at the expense of other technologies and formulates the argument that we may have invested too much in the automation applications of artificial intelligence and not enough in the uses of this new technological platform to increase human tasks and productivity. If so, the current pandemic will exacerbate these trends, with potential costs in terms of future jobs and income inequality.

Finally, the current episode may have already increased the power of tech companies, and to the extent that there were already concerns about economic concentration and the rising social power of these companies (Zuboff 2019), this is another area we need to think about.

CONCLUSION Overall, this paper is a very important contribution to both the policy debate on how to deal with the COVID-19 pandemic and to the emerging epidemiology-economics literature. Many papers will build on it, and I dearly hope that it will influence the policy debate.

Before Emmanuel Fahri's untimely death, the last sentence of this commentary was already written and it read: Given the enormity of the challenges we are facing, which go beyond containing the epidemic and its economic fallout, it is encouraging to see the best minds in our profession turn their energy to this area.

All I can add is that it is devastating for all of us that our profession will no longer benefit from one of its best and most inspiring leaders.

REFERENCES FOR THE ACEMOGLU COMMENT

Acemoglu, Daron, Victor Chernozhukov, Iván Werning, and Michael D. Whinston. Forthcoming. "Optimal Targeted Lockdowns in a Multi-Group SIR Model." *American Economic Review: Insights*.

Acemoglu, Daron, and Pascual Restrepo. 2020a. "Robots and Jobs: Evidence from US Labor Markets." *Journal of Political Economy* 128, no. 6: 2188–244.

Acemoglu, Daron, and Pascual Restrepo. 2020b. "The Wrong Kind of AI? Artificial Intelligence and the Future of Labour Demand." *Cambridge Journal of Regions, Economy and Society* 13, no. 1: 25–35.

Acemoglu, Daron, and James A. Robinson. 2012. *Why Nations Fail: The Origins of Power, Prosperity, and Poverty*. New York: Crown Business.

Alvarez, Fernando E., David Argente, and Francesco Lippi. 2020. "A Simple Planning Problem for Covid-19 Lockdown." Working Paper 26981. Cambridge, Mass.: National Bureau of Economic Research. https://www.nber.org/papers/w26981.

Atkeson, Andrew. 2020. "What Will Be the Economic Impact of Covid-19 in the US? Rough Estimates of Disease Scenarios." Working Paper 26867. Cambridge, Mass.: National Bureau of Economic Research. https://www.nber.org/papers/w26867.

Chidambaram, Priya, Rachel Garfield, and Tricia Neuman. 2020. "COVID-19 Has Claimed the Lives of 100,000 Long-Term Care Residents and Staff." Blog post, November 25, Kaiser Family Foundation Policy Watch. https://www.kff.org/policy-watch/covid-19-has-claimed-the-lives-of-100000-long-term-care-residents-and-staff/.

Chu, Derek K., Elie A. Akl, Stephanie Duda, Karla Solo, Sally Yaacoub, Holger J. Schünemann, and others. 2020. "Physical Distancing, Face Masks, and Eye Protection to Prevent Person-to-Person Transmission of SARS-CoV-2 and COVID-19: A Systematic Review and Meta-Analysis." *Lancet* 394, no. 10242: 1949–2020.

Diekmann, Odo, Johan Andre Peter Heesterbeek, and Johan A. J. Metz. 1990. "On the Definition and the Computation of the Basic Reproduction Ratio R_0 in Models for Infectious Diseases in Heterogeneous Populations." *Journal of Mathematical Biology* 28, no. 4: 365–82.

Draief, Moez, and Laurent Massoulié. 2009. *Epidemics and Rumours in Complex Networks*. Cambridge: Cambridge University Press.

Ferguson, Neil M., Daniel Laydon, Gemma Nedjati-Gilani, Natsuko Imai, Kylie Ainslie, Marc Baguelin and others. 2020. *Report 9: Impact of Non-Pharmaceutical Interventions (NPIs) to Reduce COVID-19 Mortality and Healthcare Demand*. London: Imperial College COVID-19 Response Team.

Greenhalgh, Trisha, Manuel B. Schmid, Thomas Czypionka, Dirk Bassler, and Laurence Gruer. 2020. "Face Masks for the Public during the Covid-19 Crisis." *British Medical Journal* 369, m1435. doi: 10.1136/bmj.m1435.

Hansen, Lars Peter, and Thomas J. Sargent. 2008. *Robustness*. Princeton, N.J.: Princeton University Press.

Hulten, Charles R. 1978. "Growth Accounting with Intermediate Inputs." *Review of Economic Studies* 45, no. 3: 511–18.

Jackson, Matthew O. 2010. *Social and Economic Networks*. Princeton, N.J.: Princeton University Press.

Kermack, William Ogilvy, and Anderson G. McKendrick. 1927. "A Contribution to the Mathematical Theory of Epidemics." *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 115, no. 772: 700–21.

Mossong, J., N. Hens, M. Jit, P. Beutels, K. Auranen, R. Mikolajczyk, and others. 2008. "Social Contacts and Mixing Patterns Relevant to the Spread of Infectious Diseases." *PLoS Med* 5, no. 3: e74.

Zhou, Kemin, and John C. Doyle. 1998. *Essentials of Robust Control*. Upper Saddle River, N.J.: Prentice Hall.

Zuboff, Shoshana. 2019. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. New York: Profile Books.

GENERAL DISCUSSION Justin Wolfers highlighted the paper's need for published standard errors. He noted that the exponential behavior of an epidemic should create a tremendous bias toward conservative parameter estimates.

Alan Auerbach expressed appreciation for the paper and subsequent discussion. He suggested that the paper should consider the broader feasibility of the interventions it proposes. He noted that current outbreaks are in some cases being driven by skepticism or exhaustion with social distancing and face masks. He highlighted that the practical sustainability of a policy will have an effect on the trade-offs the paper is studying. He observed that some of the policies the paper deems less attractive may in fact be more feasible than the policies it recommends.

Thomas Philippon complimented the paper and discussion. He then observed that what determines a policy's effectiveness is how behaviors change in response to the policy relative to the behavior in the absence of the policy. He highlighted the paper's proposed policy to isolate older-age cohorts and pointed out that as relatively high-risk and low-contact groups, compelling them to self-isolate creates a behavioral difference that is relatively small compared to the behavioral difference in young people—a low-risk, high-contact group—caused by limiting their social behavior. He supposed that even if the goal is ultimately to protect the old, it would be preferable to prevent young people from going to bars rather than locking down old people.

Austan Goolsbee highlighted the fact that data from China suggest a very high danger of cross-infection in the home.[1] He noted that in the United States there is a profound prevalence of COVID-19 risk factors outside of age, like obesity and other medical conditions. He suggested that this riskiness, combined with infections driven by home contact, may make targeted lockdowns unrealistic. He also suggested that using the Current Population Survey to create a matrix of spousal employment sectors would improve the model of at-home contacts.

Jim Stock responded to Philippon by explaining that the isolation driven by targeted lockdowns of the elderly is a less important behavioral determinant than changes to nursing home administration and capacity. He responded to Daron Acemoglu's strong recommendation that the paper

1. Qin-Long Jing and others, "Household Secondary Attack Rate of COVID-19 and Associated Determinants in Guangzhou, China: A Retrospective Cohort Study," *The Lancet* 20, no. 10 (2020): 1141–50, https://www.thelancet.com/journals/laninf/article/PIIS1473-3099(20)30471-0/fulltext.

adopt an optimal control framework by explaining that he and his coauthors have considered such a framework at length but have ultimately decided that there is a real value to modeling people complying with CDC guidelines in circumstances as complex as this pandemic.

Emmanuel Farhi added that in the interest of simplifying the paper's models, he and his coauthors have made a number of assumptions that mitigate the economic impact of lockdowns, as discussed briefly in the paper. He stated that they could easily have worsened the economic outlook of lockdowns by incorporating more realistic assumptions, like complementarities and Keynesian spillovers, among other things.

Farhi also highlighted that the overall message of the paper is simple: contacts in the workplace largely do not drive infections, so what will prevent a second wave are policies that reduce contacts in settings outside the home or workplace. He noted that the primary outstanding question is how effectively that goal can be achieved, and what behavioral changes will be induced versus driven by policy itself. He states that this message is very simple and is explored fairly robustly in the paper.