

1 **Adaptive policies for use of physical distancing interventions during** 2 **the COVID-19 pandemic**

3 **Reza Yaesoubi,^{1*} Joshua Havumaki,² Melanie H. Chitwood,² Nicolas A. Menzies,³ Gregg Gonsalves,²**
4 **Joshua Salomon,⁴ A David Paltiel,¹ and Ted Cohen²**

5 1. Department of Health Policy and Management, Yale School of Public Health, New Haven, CT.

6 2. Department of Epidemiology of Microbial Diseases, Yale School of Public Health, New Haven, CT.

7 3. Department of Global Health and Population, Harvard T H Chan School of Public Health, Boston,
8 MA, USA

9 4. Department of Medicine, Stanford University, Palo Alto, CA, USA

10 * Correspondence to: reza.yaesoubi@yale.edu

11

12 **One-sentence summaries:** Adaptive physical distancing policies save more lives with fewer weeks of
13 intervention than policies which prespecify the length and timing of interventions.

14 **Abstract**

15 Policymakers need decision tools to determine when to use physical distancing interventions to maximize the
16 control of COVID-19 while minimizing the economic and social costs of these interventions. We develop a
17 pragmatic decision tool to characterize adaptive policies that combine real-time surveillance data with clear
18 decision rules to guide when to trigger, continue, or stop physical distancing interventions during the current
19 pandemic. In model-based experiments, we find that adaptive policies characterized by our proposed approach
20 prevent more deaths and require a shorter overall duration of physical distancing than alternative physical
21 distancing policies. Our proposed approach can readily be extended to more complex models and interventions.

22 **Main Text**

23 The health and economic costs of the ongoing SARS-CoV-2 pandemic are staggering. In the absence of an
24 effective vaccine or other pharmaceutical interventions, physical distancing (PD) measures have been the primary
25 means to reduce the speed of epidemic growth, to relieve pressure on the health care system, and to buy scientists
26 time to develop new prevention and treatment strategies. Model-based analysis of previous pandemics (1) as well
27 as current models of the COVID-19 pandemic suggest that these PD interventions will need to be maintained for
28 weeks to months in the United States to minimize risks of overwhelming hospital capacity and to avoid epidemic
29 rebound (2, 3). However, PD measures also impose an immense economic and social burden. Indeed, it is clear
30 that a single uninterrupted period of strict PD cannot be practically sustained to achieve eradication. Accordingly,
31 policymakers require tools to help them determine when PD interventions should be started and stopped to
32 optimize health and economic outcomes.

33 Here, we describe a method for adaptive decision making during the ongoing pandemic that explicitly accounts
34 for the tradeoffs between the health benefits and economic costs of PD interventions. We demonstrate how this
35 decision model can assemble the necessary evidence to characterize policies that use surveillance data to inform

36 real-time determination of starting, continuation, or stopping of PD interventions. Using a simulation model of the
37 COVID-19 pandemic, we show that these types of adaptive policies can avert more deaths while requiring a
38 shorter total duration of PD compared to alternative policies, including pre-determined periodic interventions and
39 those that focus on ensuring available ICU capacity is not surpassed (2, 3).

40 For the analysis presented here, we assume that a policymaker's objectives are to minimize both the number of
41 deaths associated with COVID-19 and the duration of physical distancing interventions as a proxy for the harm
42 inflicted by shutdowns on the economy. We quantify the trade-off between these objectives using the net
43 monetary benefit (NMB) framework (4, 5), expressed as $\sum_{k=1}^T (\omega q_k + a_k)$, where T is the number of weeks the
44 epidemic lasts; q_k is the number of deaths due to COVID-19 that occur in week k ; and $a_k = 1$ if physical
45 distancing is in effect during week k , and $a_k = 0$, otherwise.

46 In the objective function described above, ω represents how a specific policymaker weighs the number of
47 COVID-19 deaths per 100,000 population against the duration of PD. For example, a policymaker with trade-off
48 value $\omega = 0.1$ weeks per COVID-19 death averted is willing to keep PD in place for one additional week if it
49 could prevent 10 deaths related to COVID-19 per 100,000 population. Higher values of ω indicate a higher
50 tolerance for paying the economic and social costs of PD interventions in the interest of reducing COVID-19-
51 related mortality. In the Supplement, we describe how this objective function can be extended to incorporate more
52 comprehensive outcomes including quality-adjusted life-years lost and cost incurred due to cases of COVID-19,
53 and the societal and economic cost of PD in each week.

54 Here, we estimate q_k using a simple SEIR model of SARS-CoV-2 transmission in a population of size 1,000,000.
55 Informed by data from the COVID-19 pandemic in U.S. (6), the model projects the weekly incidence of cases,
56 patients requiring hospitalization and/or critical care, and deaths due to COVID-19. We assume that the PD
57 intervention reduces the population contact rate by 70% with uncertainty interval of [60% – 80%] (3).
58 Additional details are provided in the Supplement.

59 An adaptive policy for informing PD interventions uses surveillance data available at the start of week k to
60 recommend whether PD should be started, continued, or stopped during week k . One example of such a policy
61 would recommend triggering PD when the estimated prevalence of infection surpasses 375 cases per 100,000
62 people and stopping PD when this epidemiological measure falls below 100 cases per 100,000 people (3). As data
63 such as the weekly number of COVID-19 patients requiring critical care (C_k) may be more easily observed than
64 prevalence of infection, we focus on adaptive policies that use C_k to guide the use of PD interventions during the
65 next week. We consider policies that, for a given trade-off value ω , specify two thresholds ($T_1(\omega), T_2(\omega)$) which
66 govern the recommendation of PD for the following week. This results in very simple decisions rules: If $C_k >$
67 $T_1(\omega)$, PD should be used in the following week and if $C_k < T_2(\omega)$, PD should be stopped. In the Supplement,
68 we describe the optimization algorithm which identifies the thresholds ($T_1(\omega), T_2(\omega)$) that minimize the loss in
69 the population NMB for a given trade-off value ω .

70 Fig. 1 presents an example of a policy which uses C_k to guide weekly decisions about the use of PD and
71 minimizes the loss in the population NMB. At a trade-off value $\omega = 0.1$ weeks per death averted per 100,000
72 persons, for example, this policy suggests that the PD intervention should be started when C_k exceeds a threshold
73 13.2 per 100,000 persons (Fig. 1A) and stopped when C_k falls below 12.9 per 100,000 persons (Fig. 1B). For a
74 policymaker with an uncertain trade-off value ω , Fig. 1 presents an affordability curve (Fig. 1C), which returns the
75 cumulative number of weeks for which PD interventions are expected to be used for different values of ω . A

76 policymaker with a maximum tolerance for the overall number of weeks of PD can consult Fig. 1C to select a
77 trade-off threshold that satisfies this constraint in expectation. The policymaker can then use Fig. 1A-B to make
78 real-time decisions given their selected trade-off threshold. We note that higher trade-off values lead to lower
79 thresholds of C_k at which to start PD interventions (Fig. 1A-B), more weeks during which PD is used (Fig. 1C)
80 and a lower expected number of deaths due to COVID-19 over the course of the pandemic (Fig. 1D).

81 As we have shown in previous studies of epidemics (7-9), the performance of policies to guide decision-making
82 depends on the specific features of surveillance data (e.g. C_k) selected to inform decisions. To demonstrate the
83 differential performance of PD policies in the context of the ongoing epidemic, we used our simulation model of
84 the COVID-19 pandemic to simulate the outcomes of three types of PD policies that use different features and
85 decision rules to inform recommendations (Fig. 2). We consider a *static* periodic policy (purple curve) that
86 employs the periodic use of PD at pre-defined durations and frequencies (e.g. every 2 weeks for 2 weeks, every 4
87 weeks for 4 weeks, etc.). We also evaluate two *adaptive* policies that use C_k to determine whether PD should be
88 started, continued or stopped for the following week: 1) an ‘Adaptive: ICU Capacity’ policy which is similar to
89 policies proposed by (2, 3) where the on/off PD thresholds are determined to ensure that the expected probability
90 of surpassing available ICU capacity is below a certain value (e.g. 90%); and 2) an ‘Adaptive: Minimize Loss of
91 NMB’ that determines the on/off thresholds to minimize the loss in the population NMB using the optimization
92 algorithm described in the Supplement (this policy is displayed in Fig. 1A-B).

93 The adaptive policies that seek to minimize the loss in NMB can be designed to *dominate* (i.e. prevent more
94 deaths and require shorter duration of PD as) other policies described above (Fig. 2). For any given number of
95 weeks the policy maker is willing to maintain the PD intervention, the horizontal distance between the green curve
96 and other curves in Fig. 2 represents how many deaths could be prevented by using an adaptive policy designed to
97 minimize the loss in NMB, as compared to one of the other two approaches.

98 The results presented here highlight the need for several specific data items to promote better decision-making
99 during this pandemic. First, we need the best data possible to understand the natural history of the pathogen (e.g.
100 duration of immunity, infectiousness, role of asymptomatic transmission, mortality), the current epidemiology
101 (e.g. the prevalence of infectious individuals and hospitalized patients), and the effectiveness (e.g. reduction in
102 contact rate) of PD interventions. Second, we need better data on the economic costs of physical distancing
103 interventions. Third, and most difficult, is the need for policymakers to establish what economic tradeoffs we are
104 willing to make to preserve the health of the population. In the absence of clearly established guidelines on
105 societal willingness-to-pay for health, policymakers may choose some proxies for the consumption of societal
106 resources such as the number of weeks over which PD interventions are expected to be used (Fig. 1C) to choose
107 an efficient PD policy to follow.

108 This paper presents a simple model of the COVID-19 pandemic in order to facilitate exposition and discussion.
109 As such, we have made a number of assumptions. We average across many important sources of heterogeneity.
110 Additional details (e.g. age-structure, risk-group membership, asymptomatic transmission, population compliance
111 with different PD interventions over time, availability of tests, and lags in the reporting of test results) will need to
112 be added to allow for context-specific usage and to allow for simulation of more complex interventions (e.g.
113 contact tracing, age-specific relaxation of PD rules, use of serological tests to establish which individuals should
114 not be limited by PD interventions). We used crude estimates to project health loss associated with the COVID-19
115 pandemic ignoring comorbidities and downstream health consequences of unemployment during tightened PD
116 interventions. These factors could also be incorporated, depending on the availability of data, in the objective

117 function we described above and in the Supplement. The effectiveness of different PD interventions may not be
118 known a priori, and our simple approach for modeling PD ignores the complexities in which mobility may change
119 over the course of the epidemic regardless of the PD policy adopted. New data on the relationship between
120 mobility, contact patterns, and the types of PD interventions guidelines in place can help improve the usefulness
121 of these models.

122 The decision tool we present here provides a flexible framework for using real-time observations to guide the use
123 of PD interventions. It seems likely to us that policymakers would prefer policies that are responsive to updated
124 surveillance and do not prematurely commit them to future actions. The algorithm we propose to optimize PD
125 policies does not restrict the type of epidemic model used to project the health and economic outcomes of the
126 COVID-19 pandemic. Various frameworks, including compartmental, agent-based, or network models may be
127 used as needed to relax the simplifying assumptions of our model. Furthermore, our proposed approach can also
128 incorporate delays in reporting of data and uncertainties in estimates (due to limited sample size) (7, 9). We have
129 described an algorithm to optimize PD policies in the Supplement and welcome other groups to contact us or
130 adapt this approach to help improve real-time decision making during this epidemic.

131 **References**

- 132 1. M. C. Bootsma, N. M. Ferguson, The effect of public health measures on the 1918 influenza pandemic in
133 U.S. cities. *Proc Natl Acad Sci U S A* **104**, 7588-7593 (2007).
- 134 2. N. M. Ferguson *et al.*, "Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality
135 and healthcare demand," (Imperial College, 2020).
- 136 3. S. M. Kissler, C. Tedijanto, E. Goldstein, Y. H. Grad, M. Lipsitch, Projecting the transmission dynamics
137 of SARS-CoV-2 through the postpandemic period. *Science*, (2020).
- 138 4. G. D. Sanders *et al.*, Recommendations for Conduct, Methodological Practices, and Reporting of Cost-
139 effectiveness Analyses: Second Panel on Cost-Effectiveness in Health and Medicine. *JAMA* **316**, 1093-
140 1103 (2016).
- 141 5. A. A. Stinnett, J. Mullahy, Net health benefits: a new framework for the analysis of uncertainty in cost-
142 effectiveness analysis. *Med Decis Making* **18**, S68-80 (1998).
- 143 6. C. C.-R. Team, Severe Outcomes Among Patients with Coronavirus Disease 2019 (COVID-19) - United
144 States, February 12-March 16, 2020. *MMWR Morb Mortal Wkly Rep* **69**, 343-346 (2020).
- 145 7. R. Yaesoubi, T. Cohen, Identifying dynamic tuberculosis case-finding policies for HIV/TB coepidemics.
146 *Proc Natl Acad Sci U S A* **110**, 9457-9462 (2013).
- 147 8. R. Yaesoubi, T. Cohen, Identifying cost-effective dynamic policies to control epidemics. *Stat Med* **35**,
148 5189-5209 (2016).
- 149 9. R. Yaesoubi *et al.*, Adaptive guidelines for the treatment of gonorrhea to increase the effective life span of
150 antibiotics among men who have sex with men in the United States: A mathematical modeling study.
151 *PLoS Med* **17**, e1003077 (2020).

152

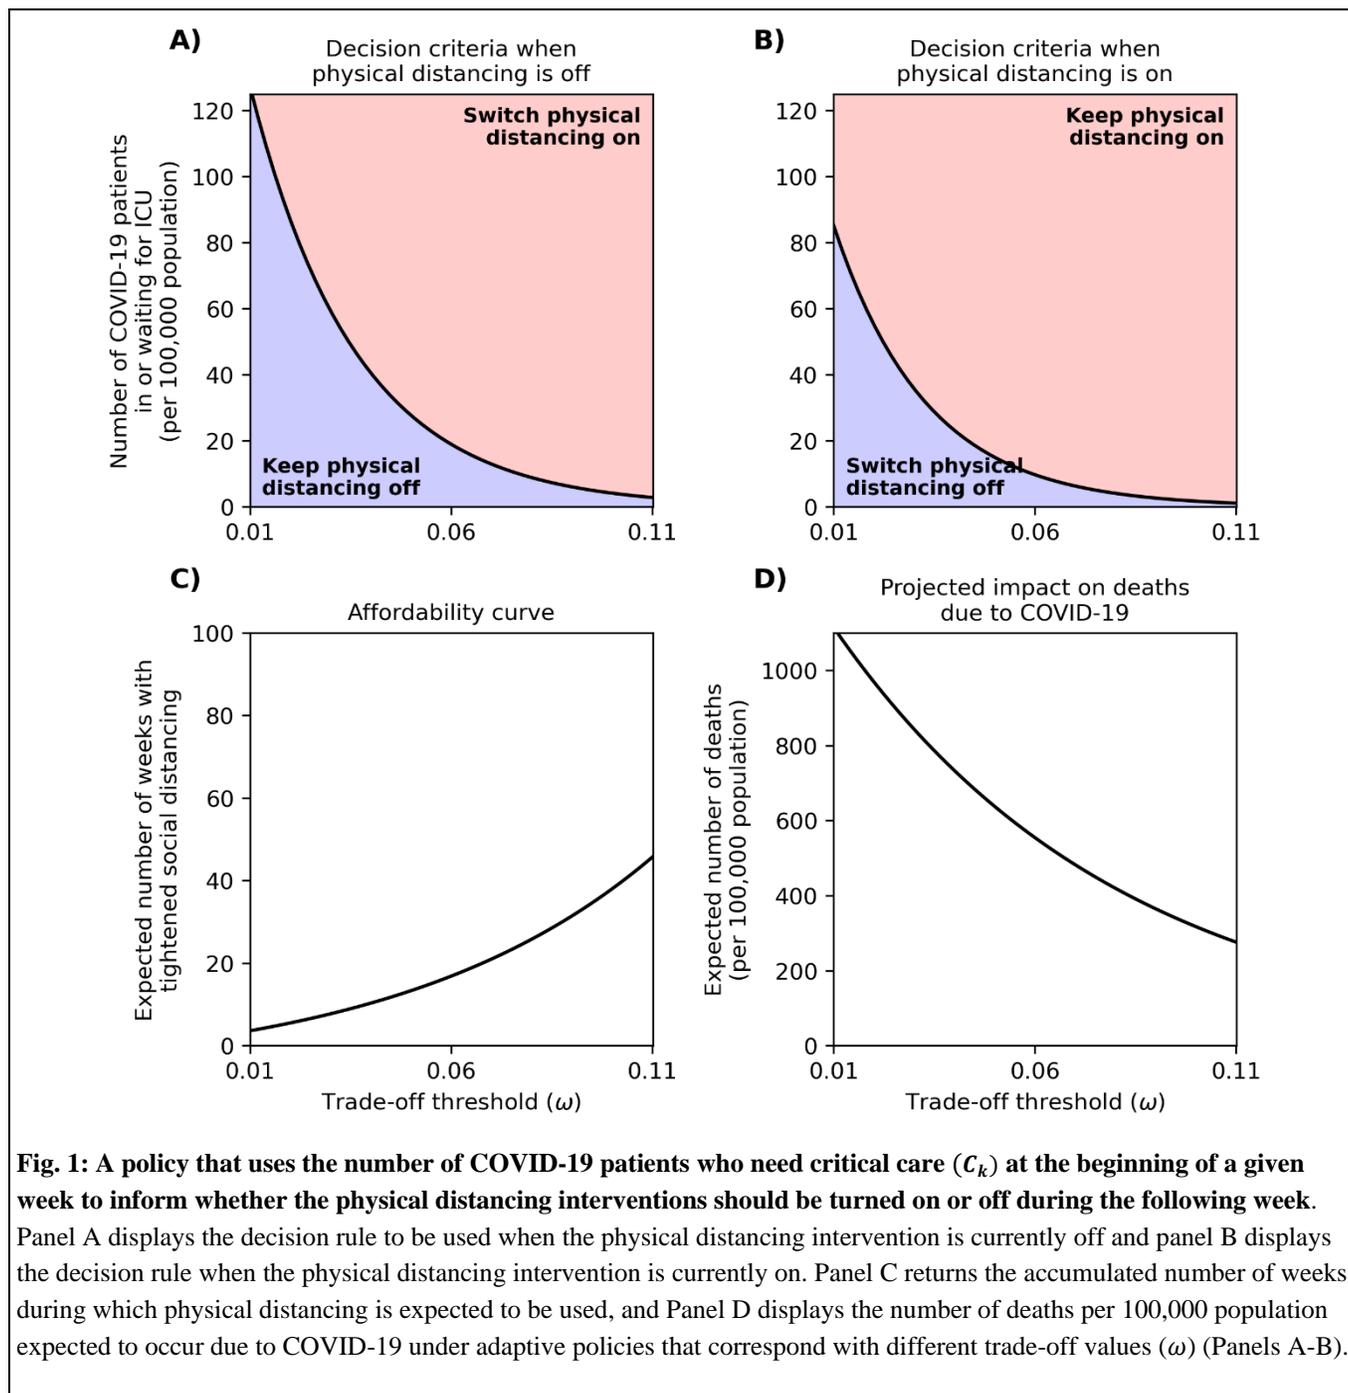
153 **Acknowledgments:** RY was supported by K01AI119603, and TC by R01AI112438 and R01AI146555, all from
154 the National Institute of Allergy and Infectious Diseases. GG was supported by DP2DA49282 and ADP by
155 R37DA015612 both from the National Institute on Drug Abuse. Author contributions: RY: Conceptualization,
156 Formal Analysis, Investigation, Methodology, Software, Visualization, Writing (original draft); JH: Data curation,
157 Writing (review and editing); MC: Data curation, Writing (review and editing); NAM: Supervision, Writing
158 (review and editing); GG: Supervision, Writing (review and editing); JS: Supervision, Writing (review and

159 editing); ADP: Supervision, Writing (review and editing); TC: Conceptualization, Funding acquisition,
160 Methodology, Supervision, Writing (original draft). Authors declare no competing interests. All data is available
161 in the main text or the supplementary materials. The code of the model is provided at
162 <https://github.com/yaesoulab/APACE>.

163 **Supplementary Materials:** Materials and Methods, Figures S1, and Tables S1-S2.

164

165 **Figures**



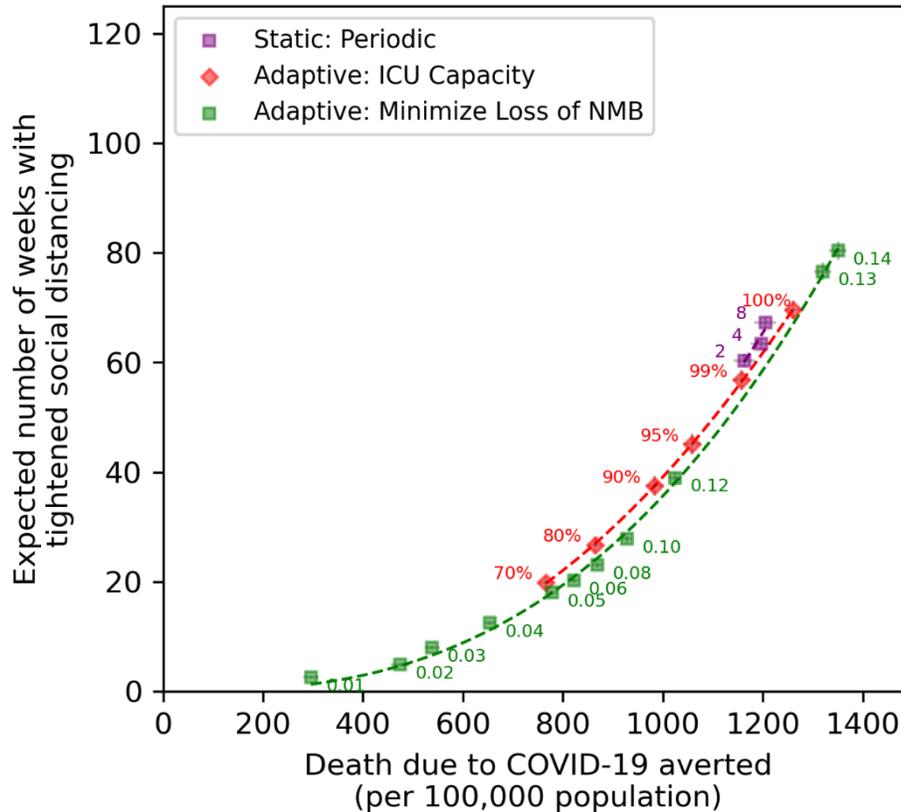


Fig. 2: Comparing the performance of difference physical distancing (PD) policies. The origin in this figure corresponds to a counterfactual scenario where PD is not used at all, which is expected to result in 1,391 deaths per 100,000 population. The curve ‘Static: Periodic’ corresponds to policies that recommend using PD every 2 weeks for 2 weeks, every 4 weeks for 4 weeks, and so on. The numbers on the red curve (‘Adaptive: ICU Capacity’) represent the probability that a COVID-19 patient who requires critical care would get admitted to ICU. The numbers on the green curve (‘Adaptive: Minimize Loss of NMB’) represent the trade-off thresholds of policies selected from Fig. 1A-B. The bars represent the 95% confidence interval of projections using 200 simulated trajectories. NMB: net monetary benefit.