Review of Economic Design
https://doi.org/10.1007/s10058-025-00379-z

ORIGINAL PAPER

$R < 1$ as an economic constraint

Eric Budish[1]

Received: 20 June 2023 / Accepted: 12 February 2025

Abstract
This paper proposes a novel pandemic response paradigm, and shows that it would have been the right middle ground between lockdown and ignore-the-virus for Covid-19: maximize social welfare subject to $R \leq 1$ as a *constraint*. A simple graphical argument shows that this formulation is an approximately optimal way to balance socioeconomic and health objectives, because of a sharp kink in the benefits-of-risk-reduction curve at $R = 1$ (both the curve and its kink are novel to this paper). Two critical insights emerge from this approach to the pandemic. First, the $R \leq 1$ constraint imposes a "risk budget" on society. Society should optimally spend this budget on the social and economic activities with the highest ratio of socioeconomic value to disease-transmission risk, with targeted activity bans for activities with too low a ratio of value-to-risk. For example, schools have a much higher ratio of value-to-risk than bars, so society should optimally spend its risk budget on schools over bars. Second, what I call "low-cost risk reducers" (LCRRs) can significantly improve activities' value-to-risk ratios and hence significantly reduce the cost of satisfying the $R \leq 1$ constraint. Examples of LCRRs for Covid-19 include rapid testing, high-quality facemasks, stay-home-if-sick rules and improved air circulation. A simple numerical example, based on estimates from the medical literature for R_0 and the efficacy of LCRRs for Covid-19, suggests the potential gains from this paper's approach to the pandemic would have been enormous—plausibly trillions of dollars and hundreds of thousands of lives in the United States alone.

An earlier version of the ideas in this paper circulated on April 1, 2020 (Budish 2020a). The April 2020 draft and conference presentations on this work from April and May 2020 are available on the author's website. Special thanks are due to Chris Avery and David McAdams for co-editing this special issue and for valuable comments and advice. I also thank Jason Abaluck, Nikhil Agarwal, Mohammad Akbarpour, Matteo Aquilina, Emma Berndt, Judy Chevalier, John Cochrane, Michael Droste, Amy Finkelstein, Joshua Gans, Austan Goolsbee, Sam Hanson, Ralph Koijen, Scott Kominers, Michael Lewis, Shengwu Li, Yueran Ma, Neale Mahoney, Emi Nakamura, Emily Oster, James Stock, Adi Sunderam, Chad Syverson, Alex Tabarrok, Sarah Taubman, Heidi Williams and an anonymous referee for valuable discussions and feedback at various stages of this research. I thank Ethan Che, Jiahao Chen, Jia Wan, Zizhe Xia, and Tianyi Zhang for research assistance. I have no financial interests or conflicts relevant to this research.

This paper is dedicated to my partners in lockdown and life, Emma, Nathan and Jacob.

✉ Eric Budish
eric.budish@chicagobooth.edu

[1] University of Chicago Booth School of Business, Chicago, IL, USA

Published online: 12 May 2025

Keywords Pandemics · Covid-19 · $R < 1$ · Pandemic playbook

JEL Codes C61 · D47 · D6 · D7 · E6 · H0 · H12 · I1 · I18 · L51

1 Introduction

Economists usually think about public policy with a constrained maximization problem in the back of their heads, something like

max Social Welfare (1)
subject to
Technological Constraints
Incentive Constraints

where "Social Welfare" includes both the economic and non-economic dimensions of well-being. Traditional economic constraints include technology (we can't all have infinite of everything) and incentives (it is hard to get someone to work hard or innovate without compensation).

Health policy experts, in their response to the Covid-19 crisis, often seemed to have a constrained maximization problem in the back of their heads like

min Spread of Covid-19 (2)
subject to
Keeping Society Functioning

This perspective is understandable as a response to the fear of exponential growth of a deadly threat. Prior to intervention, Covid-19 cases doubled every few days. Left unchecked, such exponential growth would quickly overwhelm medical systems and lead to tens of millions of deaths globally.

The difficulty with (2), however, is that it makes it impossible to think about trade-offs. The extreme versions of social distancing that are called for by (2)—closing schools, shuttering entire industries, avoiding close interactions with other people—themselves have enormous costs. The March 2020 Imperial College epidemiological model, which reportedly influenced lockdown decisions in many countries, discussed the possibility of children's schools being closed for up to 2 years.[1] Lockdowns conservatively cost the global economy $1 trillion per month (Castillo et al. 2021). But formulation (2) does not allow for *any* unnecessary risk. Dr. Francis Collins, former head of the National Institutes of Health, said in Summer 2023 of the "public-health mindset":

> "If you're a public-health person and you're trying to make a decision, you have this very narrow view of what the right decision is, and that is something that will save a life. Doesn't matter what else happens. So you attach *infinite value*

[1] Source: Imperial College Covid-19 Response Team, "Impact of non-pharmaceutical interventions (NPIs) to reduce COVID19 mortality and healthcare demand," March 16, 2020. Excerpt from page 11: "The right panel of Table 4 shows that social distancing (plus school and university closure, if used) need to be in force for the majority of the 2 years of the simulation."

> *to stopping the disease* and saving a life. You attach a *zero value to whether this actually totally disrupts people's lives, ruins the economy, and has many kids kept out of school in a way that they never quite recovered. Collateral damage.* This is a public-health mindset. And I think a lot of us involved in trying to make those recommendations had that mindset and that was really unfortunate. That's another mistake we made." (Collins 2023. Emphasis added)

Formulation (2) also invariably leads to political fights over what it takes to keep society functioning without a clear decision framework. Fully 40% of the US adult population was designated an "essential worker" (McCormack et al. 2020).

This paper proposes a novel pandemic response paradigm that incorporates a version of the pure health perspective (2) into the traditional economic perspective (1), and shows that it would have been an approximately optimal way to balance traditional social and economic objectives with health objectives. My approach is also simple and intuitive, using plain vanilla static optimization (as opposed to a more complicated and opaque dynamic model, to which I will return below), and focuses attention on what I will argue is the right set of policy issues. My proposed paradigm is:

max Social Welfare (3)
  subject to
Technological Constraints
Incentive Constraints
$R \leq 1$ Constraint: Reduce the Covid-19 Average Transmission Rate to Below 1
(Until Vaccines or Treatments are Widely Available)

This formulation superficially looks like (1)—in particular, it has the usual economic objective function, social welfare, which is the polar opposite of placing "zero value" on disrupting people's lives, ruining the economy, and keeping kids out of school as in the Dr. Francis Collins quote above. But, operationally, it will also do well at approximating the health objective in (2), because of the additional constraint that has been added: reduce "R", the average transmission rate of Covid-19, to below the critical threshold of 1.

As is widely understood, diseases with average transmission rates above 1 eventually infect huge numbers of people, whereas diseases with average transmission rates below 1 do not. In the case of Covid-19 in particular, with both a particularly rapid spread on the order of a few days to a week, and the perceived likelihood of a successful vaccine within a year or two (which turned out to be correct), the difference between $R > 1$ and $R \leq 1$ is particularly dramatic. For example, a reproductive rate of $R = 1.30$ would translate to 140 million infections and 979,000 deaths in the United States in 12 months, whereas a reproductive rate of $R = 1.00$ would translate to just 5.8 million infections and 40,000 deaths in that same time period (each assuming an initial stock of 100,000 infections, a 5-day infectiousness period, and an infection fatality rate of 0.7%; see Fig. 1 and Table 1). If society engineers R less than 1 then the scale of the health crisis, while not *absolutely* minimized as in formulation (2), is *approximately* minimized. At the same time, by having traditional economic and societal goals as the objective function, (3) leads to very different policy implications. In particular, formulation (2) inevitably leads to an as-severe-as-feasible societal lockdown, whereas formulation (3) seeks to allow as much socially-valuable activity as possible subject to $R \leq 1$.

I start the analysis by showing why formulation (3) is approximately optimal. The key point is that the benefits of reducing disease transmission sharply increase as R approaches $R = 1$ from above, but then flatline beyond this level of risk reduction. That is, there is a "kink" in what I call the benefits-of-risk-reduction curve (both this curve and its kink are novel to this paper). At the same time, the cost of reducing disease transmission has the standard convex-increasing shape, reflecting that reducing risk becomes increasingly expensive as society goes from the easy risk reductions to the more expensive risk reductions.[2] The non-standard benefits curve with the kink at $R = 1$ and the standard convex-increasing cost curve combine to imply that $R \leq 1$ is approximately the optimal policy target for cost and benefit parameters that seem reasonable for the case of Covid-19. See Sect. 2 and especially Fig. 2 for the overall conceptual argument, Proposition 3 in Sect. 3 for a formal approximate optimality claim, and see Fig. 7 later on in Sect. 7 for a version of the optimality argument calibrated to realistic parameters for Covid-19.

I then analyze two critical insights that emerge from formulating society's optimization problem with $R \leq 1$ as a constraint. First, the $R \leq 1$ constraint imposes a "*risk budget*" on society. Society should optimally spend its risk budget on the social and economic activities that maximize the ratio of socioeconomic value to disease-transmission risk—in effect, an activity's social welfare "bang for buck" per unit of virus risk. This means that some activities should optimally be allowed despite having relatively high risk, and vice versa. In math, this bang for buck is given by $\frac{v}{r}$, where v is an activity's social value (benefits less costs) and r is its contribution to the spread of the virus. I call banning an activity with low $\frac{v}{r}$ a "*targeted activity ban*." This analysis is presented in Sect. 3.

Second, society should try to satisfy the $R \leq 1$ constraint as cheaply as possible. What I will call "*low-cost risk reducers*" (LCRRs)—meant to encompass rapid testing, high-quality facemasks, improved air circulation, six feet of social distance, and other related ideas—can be conceptualized as significantly reducing the transmission risk of activities, i.e., their risk r, at low cost to their value v relative to not having the activity at all.[3] Such interventions thus significantly improve activities' $\frac{v}{r}$ ratios, which in turn expands the production possibilities frontier of how much social welfare can be achieved while keeping within the constraint $R \leq 1$. Said differently, LCRRs decrease the societal cost of satisfying the $R \leq 1$ constraint. Section 4 models these ideas formally. Section 5 presents a detailed numerical example, using estimates from the medical literature for R_0 and the efficacy of LCRR's, and uniform distributions for value and risk.

[2] Many dynamic models in economics had just a single representative activity (e.g., "working" or "not working"), whereas my model has heterogeneous activities that vary in both value and risk. Thus, the cost-of-risk-reduction curves implicit in these dynamic models did not reflect that some risk reductions are much easier and cheaper than others. This is one reason these models missed my focus on $R \leq 1$ as a feature of the optimum. The other reason is that they did not allow for low-cost risk reducers in a meaningful way. See further discussion below.

[3] I use the phrase LCRRs instead of Non-Pharmaceutical Interventions (NPIs) because the term NPIs as used in the public health literature encompasses interventions that are both low cost and very high cost, like severe lockdowns and the closure of schools (Ferguson et al. 2020).

The numerical example highlights just how valuable LCRRs can be. If $R_0 = 2.5$ and activities' social value and risk are uniformly distributed, then without LCRRs society has to drop fully 45% of activities, costing 27% of social welfare, to get to $R \leq 1$. This is a severe societal lockdown. Whereas if LCRRs can reduce transmission risk by 50%—which is a plausible magnitude for either rapid tests or high-quality facemasks on their own and conservative if the full suite of LCRRs is deployed—then society can maintain 85% of its pre-virus activities and 97% of its pre-virus social welfare (gross of the cost of the LCRRs), while still satisfying $R \leq 1$. If some activities are "super spreaders" then the math can be even more attractive, because shutting down the super-spreader activities can be enough.

Section 6 briefly discusses dynamic considerations. In particular, it is very valuable to implement the $R \leq 1$ constraint early, before the stock of infections grows too large. If the stock of infections is high enough, it can be optimal to first invest in reducing the stock of infections with $R \ll 1$, and then transition to $R \leq 1$.

Section 7 then provides an overall sense of magnitudes for the gains from this paper's approach to the pandemic relative to what actually happened. The gains are plausibly trillions of dollars, hundreds of thousands of lives, tens of millions fewer pre-vaccine infections, and reducing immeasurable harm to the education and mental health of a generation of young people.

Remark 1 Schools vs. Bars
Likely one of the top few policy mistakes in response to Covid-19 was the long-term closing of schools in many countries, which significantly negatively impacted large numbers of children (Jack et al. 2023, Jack and Oster 2023, Kofoed et al. 2024, Goldhaber et al. 2023, Hanushek and Woessmann 2020, Azevedo et al. 2021). In the language of this paper's model, education has a very high social value v, and, with the right LCRR's in place, relatively low r (Oster 2020, Varma et al. 2021, Varma et al. 2022). At the same time, many cities in the United States allowed for bars to reopen before public schools. Criticism of these policies was widespread and can be understood as bars having significantly lower $\frac{v}{r}$ than education, so it is very inefficient to spend a scarce risk budget on opening bars over opening schools.

Remark 2 Accelerating Vaccine Development and Availability
Since imposing an $R \leq 1$ constraint is costly (i.e., the cost of using LCRRs and the cost of targeted activity bans), there is large social value to accelerating the development and availability of vaccines and effective treatments. In companion work with a large set of collaborators (Ahuja et al. (2021), Castillo et al. (2021)) we investigate the optimal investment in vaccine capacity "at risk", i.e., investing before it is known which vaccines will be successful, and give a sense of magnitudes for the social and economic value of accelerating vaccination. The numbers are very large. If our model in Ahuja et al. (2021) were followed, vaccination would have been completed in the United States by March 2021 and globally by October 2021 (this is without any speedups to the FDA approval process, which could have further sped up availability). The analysis in Castillo et al. (2021) suggests this acceleration would have been worth trillions of dollars and saved millions of lives.

In the conclusion of this paper I will argue that a new play in the pandemic playbook was called for by Covid-19: (i) pre-vaccine, treat $R \leq 1$ as a constraint and maximize socioeconomic welfare subject to this constraint, (ii) use LCRRs and targeted activity bans to get to $R \leq 1$ as cheaply as possible, and (iii) accelerate vaccination essentially as much as is feasible. This paper's emphasis is (i) and (ii) whereas my companion work Castillo et al. (2021) and Ahuja et al. (2021) focuses on (iii).

If there is a pandemic for which it is known that a vaccine or effective treatment is impossible, then the optimal policy might be quite different. In particular, a "flatten the curve" style policy that gets to herd immunity at least cost may be best. See Rachel (2024) for a careful analysis of this scenario, and see the conclusion of this paper for more discussion of the overall pandemic playbook.

Remark 3 Static Approximation to the Full Dynamic Model, and what the Dynamic Models Missed

This paper uses a static optimization model with traditional socioeconomic goals as the objective and $R \leq 1$ as a non-standard constraint. This constraint in turn forces the optimization to approximate health objectives.[4] While my approach is non-standard, it is simple and intuitive, and I show that it approximates the optimal policy in a more standard formulation of the problem. This is because of the combination of the convex-increasing cost-of-risk-reduction curve and the unusual benefits-of-risk-reduction curve that is convex-increasing and then has a sharp kink at $R = 1$. (See Figs. 2 and 7 and Proposition 3). My approach then focuses attention on what is the cheapest way to satisfy the $R \leq 1$ constraint, which in turn leads to analysis of targeted activity bans and LCRRs.

A more popular approach within economics has been to study dynamic optimization models with both socioeconomic goals and health goals in the objective function and with SIR disease dynamics added to the model as an additional set of dynamic constraints. In principle, this approach is more complete and intellectually satisfying than the static optimization approach that I take, at the cost of added complexity. However, the early dynamic models missed two key features of my model which in turn caused them to miss this paper's key insights and, in my view, focus on the wrong set of policy issues.[5]

First, for tractability, they assumed just a single representative activity, whereas my model has heterogeneous activities that vary in both value and risk.[6] This implicitly assumes away the possibility of targeted activity bans, which increases the cost of risk reduction implicit in the model because there is no scope for doing the easy risk reductions first before harder and harder risk reductions.

Second, the early dynamic models either did not have LCRRs such as masks and tests or only included them in a limited way. This further prevented it from being possible to get to $R \leq 1$ cheaply, where by cheap I mean in relation to the massive

[4] Non-standard constraints, approximations, and an engineering approach are more common in the economic field of market design. See Roth's (2002) famous manifesto "The Economist as Engineer" as well as my own discussion of the do's and don'ts of non-standard objectives and constraints in Budish (2012).

[5] Alvarez et al. (2020) first circulated March 23, 2020. Other early dynamic models that first circulated in March-April 2020 and made similar assumptions that precluded getting to $R \leq 1$ relatively cheaply include Acemoglu et al. (2020), Eichenbaum et al. (2020), Farboodi et al. (2020), and Jones et al. (2020).

[6] For example, in the main model of Alvarez et al. (2020), agents are either (i) "in lockdown" and not productive, or (ii) "not in lockdown," produce a homogeneous output w, and have just as many contacts and just as much infectiousness as in a society that is completely unaware of the virus.

costs of widespread lockdown or the massive costs of the virus spreading unchecked. Intuitively, in the simplest model with a single homogeneous activity and no LCRRs, reducing R from $R_0 = 2.5$ to $R \leq 1$ would take a 60% reduction of GDP, because $\frac{2.5-1.0}{2.5} = 0.60$.

Together, these two differences from my model—no scope for targeted activity bans and no or limited LCRRs—prevented $R \leq 1$ from being a feature of the optimal solution. This in turn generated very interesting, but in my view wrong, dynamics. For example, the optimal policy in Alvarez et al. (2020) does not impose any restrictions for nearly three weeks, then escalates to a lockdown rate of nearly 70% over the next few weeks, and then gradually lowers the lockdown rate until the 20th week (Fig. 1, Panel "Lockdown Policy"). In Acemoglu et al. (2020), the optimal policy features over a year of maximal lockdown for vulnerable groups ("old" in the model), while for less vulnerable groups ("young" and "middle") the level of lockdown is first zero, then gradually increases to between 25 and 70% depending on parameters and objectives, and then very gradually decreases back to zero (Figures 5.4-5.6, Panel "Lockdown Policy").

A good opportunity for future research would be to develop a fully dynamic optimization model but with (i) heterogeneous activities and scope for targeted activity bans, (ii) larger scope for LCRRs, and also (iii) some complexity cost of excessive dynamic fine tuning (see next remark). I conjecture that in such a model the optimal dynamics would become a lot simpler: just get to $R \leq 1$ cheaply and stay there until a vaccine is available. Such an analysis might also yield a more theoretically complete version of the pandemic playbook that I describe in this paper's conclusion.

Remark 4 Pandemic Fatigue

One last advantage of my static approach is that it encourages policy makers to pick a single policy principle and stick with it. Dynamic models with dynamic control variables (such as the degree of lockdown) assume a degree of policy control that seems both unrealistic and likely to exacerbate pandemic fatigue, with messages and rules changing frequently.

It is impossible to run the counterfactual, but imagine if public-health officials and policy makers had articulated early on that (i) their guiding principle was to allow for as much social and economic activity as possible while preventing exponential growth of the virus; that (ii) masks and tests are a way of keeping schools and most of the economy open; and that (iii) they would do all they could to speed up the availability of vaccines. If only!

2 Why $R \leq 1$?

2.1 Importance of $R \leq 1$ in the SIR Model.

Figure 1 illustrates the importance of the $R \leq 1$ threshold in the standard SIR epidemiological model.[7] Each line represents a different number of initial infections, ranging from 1,000 to 1 million. The horizontal axis of each panel varies R_0 ("R-nought"), the initial average transmission rate.[8] More precisely, what varies along the horizontal axis is the SIR model's β parameter (which represents the rate of infectiousness), with the γ parameter (where $\frac{1}{\gamma}$ represents the duration of infectiousness) held fixed, and with R_0 defined according to $R_0 \equiv \frac{\beta}{\gamma}$. The vertical axis then depicts the *cumulative* number of infections and deaths over a 12 month period, using a relatively conservative infection fatality rate of 0.7%.[9] That is, rather than the typical focus on the dynamic path of the virus over time (e.g., in the famous "flatten the curve" graphics, which I will show were misguided), this figure just plots the aggregate number of infections and deaths in a year.

Focus first on the middle and right of the figure. What this shows is that if the transmission rate is anywhere close to the estimates for Covid's reproductive rate without any intervention, e.g., R_0 in the roughly 2.0−4.0 range, there would be in excess of 250 million infections and 1.8 million deaths in the United States in a 12 month period, essentially irrespective of the initial seed of infections. This is because of rapid exponential growth. Now look in the range $R_0 = 1.2 - 1.5$, which corresponds to the idea of "flatten the curve." There are still 100 to 200 million infections and 700,000 to 1.3 million deaths, again, irrespective of the initial seed. Now look right around $R_0 \approx 1$. For the bottom 3 lines—initial seeds of 1,000 to 100,000, which may represent the stock of infections in the United States as of late February or early March 2020[10]—if $R_0 \approx 1$ then the cumulative number of infections is hard to discern from zero on the figure. So there is a huge payoff from reducing R to 1.0 even versus 1.2.

[7] See Avery et al. (2020) and McAdams (2021) for excellent and complementary surveys of the standard SIR model and its many variations, as well as open questions for economists.

[8] A note on notation: throughout this paper I will mostly use the notation "R", without any subscripts or arguments, to refer to the average transmission rate of Covid-19 at a moment in time as a function of any interventions, behavioral changes, or accumulating herd immunity. This is sometimes called R_e (e for "effective"), R_t (with t for evolution over time), or $\hat{R}$. I reserve the notation R_0 ("R-nought") to refer specifically to either (i) the initial transmission rate of Covid-19, or (ii) the basic reproduction number in the standard SIR model, defined as $R_0 \equiv \frac{\beta}{\gamma}$.

[9] The 0.7% figure is based on the CDC's Pandemic Planning Scenarios Current Best Estimate in Sept 2020 (Centers for Disease Control and Prevention 2020). The CDC estimate provides IFR estimates by age group which I turn into an overall IFR using Census Bureau data for population by age group. The CDC's optimistic scenario generates an IFR of 0.4% and its pessimistic scenario generates an IFR of 1.3%. The influential Imperial College modeling team used an IFR of 0.9% in Ferguson et al. (2020) and has estimated the IFR to be in the range 0.9%−1.26% for Europe (Flaxman et al. 2020). Economics papers such as Fernández-Villaverde and Jones (2022) and Stock (2020) emphasize the econometric difficulty of identifying the true number of underlying infections and hence the IFR.

[10] The first confirmed case of community spread in the United States was announced on Feb 27th. In the first week of March, there were about 250 confirmed cases. In the week leading to March 15th, there were about 3000 confirmed cases. It is widely known that actual cases meaningfully exceed reported cases, especially early in the crisis when testing was especially poor. See Stock (2020) and Stock et al. (2020).

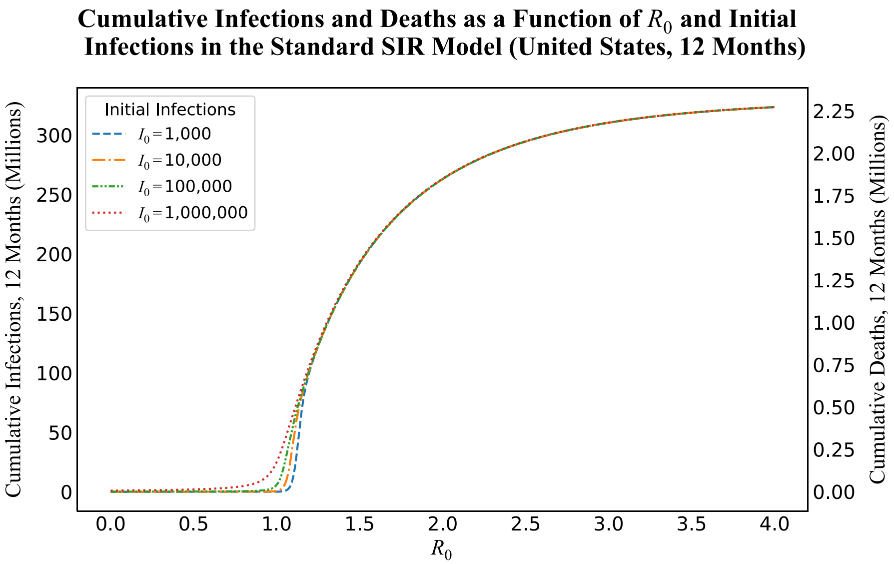

Fig. 1 *Note:* Output is based on the standard SIR model. Each line depicts a different initial infection seed. The γ parameter is fixed throughout at $\frac{1}{5}$, which represents a duration of infectiousness of 5 days. (The figure is similar if other reasonable values of γ are used instead). The β parameter, which represents the rate of infectiousness, is varied such that $R_0 = \frac{\beta}{\gamma}$ is the value depicted along the horizontal axis. The vertical axis depicts the cumulative number of infections and deaths in the United States over a 12-month period as a function of $R_0 = \frac{\beta}{\gamma}$, based on an infection fatality rate of 0.7% (per CDC estimates) and a population of 330 million

Table 1 Cumulative Infections and Deaths as a Function of R_0 (Initial Seed 100,000 Infections, 12 Months)

	R_0	# of Total Infections	# of Total Deaths
Lockdown	0.50	200,000	1,000
	0.60	250,000	2,000
	0.70	333,000	2,000
	0.80	498,000	3,000
$R \leq 1$ Approach	0.95	1,860,000	13,000
	1.00	5,810,000	40,000
Flatten the Curve	1.20	104,000,000	727,000
	1.30	140,000,000	979,000
	1.40	169,000,000	1,180,000
	1.50	192,000,000	1,350,000
Ignore	2.00	263,000,000	1,840,000
	2.50	295,000,000	2,060,000

See Notes for Fig. 1

Last, look all the way to the left of the figure, which represents a severe lockdown. What this shows is that if R_0 is low enough (e.g., $R_0 = 0.5$), the virus is essentially stopped in its tracks, and there is not much growth beyond the initial infection seed.

Table 1 complements Fig. 1 by showing the number of infections and deaths for an initial seed of $I_0 = 100,000$. Again, what leaps out from the page is the huge value of $R \leq 1$ relative to R even modestly greater than 1. Actual policy in the United States vacillated between strict lockdowns and relatively few restrictions, yielding the same outcome in the first 12 months in terms of cumulative deaths as if $R \approx 1.12 - 1.14$ (see discussion in Sect. 7).

2.2 Why R ≤ 1 is Approximately Optimal.

Figure 2 depicts a simple graphical argument for why $R \leq 1$ is likely the optimal target for economic policy. The blue line, labeled "Benefits of Risk Reduction", depicts the same information as in Fig. 1 but with both axes flipped: on the horizontal axis, further to the right now means lower R as opposed to higher, and on the vertical axis, higher now represents the number of people not infected or dead as opposed to the number infected or dead. The red line, labeled "Cost of Risk Reduction", represents the economic cost of reducing the spread of the virus, e.g., by reducing economic activity or utilizing LCRRs. The theory in Sects. 3 and 4 will microfound that this cost curve is increasing and convex and the simulation in Sect. 5 yields the specific shape used in the figure. Section 6 will discuss evidence in Goolsbee and Syverson (2021) and behavioral SIR models that suggest that the economic cost of reducing spread is potentially decreasing in the region to the left of $R = 1.0$, because of the economic damage caused by fear of contracting the virus if the stock of infections grows large, which will occur if $R > 1$. These points only enhance the case for $R \leq 1$ as an optimal policy target. The vertical scales of both the benefits and costs curves will be calibrated in Sect. 7; see Fig. 7.

Optimal policy maximizes the difference of benefits less costs. The reason that $R \leq 1$ is approximately optimal is that the benefits curve has a "kink" at $R = 1$—the health benefits of lowering R increase at an increasing rate until $R = 1$, at which point the curve not only becomes concave but essentially flat.[11] Therefore, a policy that targets $R \leq 1$ reaps almost all of the health benefits of mitigation—and, particularly the steeply increasing part as R approaches 1 from above, representing the gain from avoiding exponential growth—without incurring further, increasing, economic mitigation costs to go even further into the concave and essentially flat part of the health benefits curve.

For a formal mathematical claim that $R \leq 1$ is approximately optimal, see Proposition 3 in the next section.

[11] The figure depicts the blue benefits curve for the case of $I_0 = 100,000$. The kink is even more stark for the cases of $I_0 = 10,000$, and $I_0 = 1,000$. The transition from the convex part of the curve to the concave part is more gradual for the case of $I_0 = 1$ million, i.e., for a large-enough current stock of infections. In this case, it may be socially optimal to aim for R meaningfully less than 1, or to adopt a dynamic policy in which R is at first meaningfully less than 1, and then gradually constraints are relaxed. See discussion of these issues in Sect. 6. Farboodi et al. (2021) point out that technically the mathematically optimal choice of R might be slightly larger than 1 if the stock of infections is small enough, the confidence that a vaccine will arrive is high enough, and the policy maker has an ability to dynamically fine tune the level of lockdown.

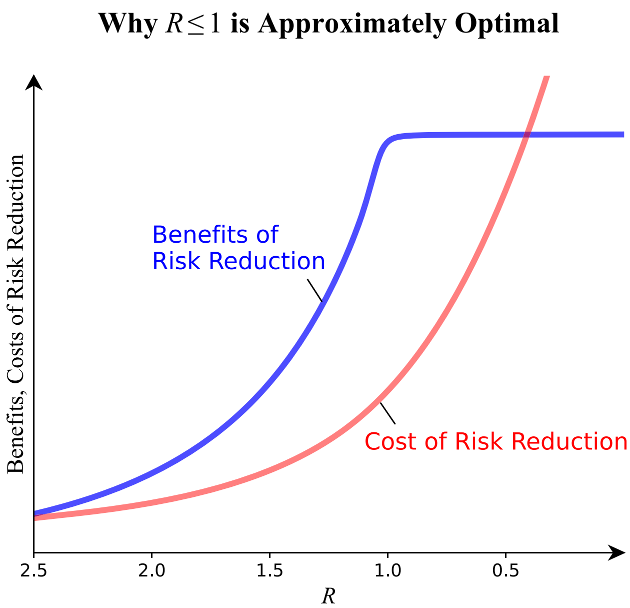

Fig. 2 *Note:* The blue line depicts the same information as the $I_0 = 100,000$ case of Fig. 1, but with both axes flipped as described in the text. The red line depicts a convex and increasing cost curve whose specific shape is based on the simulation example in Sect. 5. Both curves are depicted under the assumption that R_0 without any interventions or behavior changes is 2.5, per the CDC's best estimate as of September 2020. The vertical scales of both curves is calibrated in Sect. 7, see Fig. 7

3 Initial model (without LCRRs)

Society chooses a vector of activities $x \in X = [0, 1]^n$. Each activity i has traditional socioeconomic benefits and costs, denoted b_i and c_i, with $v_i = b_i - c_i$ the net *socioeconomic value* of the activity. Each activity i also has a *disease-transmission risk* denoted r_i. Initially, the v_i's ignore the existence of the virus; that is, the value of an activity represents its benefits less costs, including both social and economic dimensions, in the world pre-Covid 19 (see Akbarpour et al. 2024). For this reason, assume $b_i > c_i$ for all i and hence $v_i > 0$ for all i.[12] The disease-transmission risk r_i represents the activity's expected contribution to transmission of the virus in a society that does not engage in any risk reduction. For simplicity, benefits, costs, and risk are each linear in activities. Formally, activity vector $x \in X$ yields traditional socioeconomic value of $\sum_i x_i v_i$ and an effective reproduction rate of the virus of $\sum_i x_i r_i$.[13]

12 I will not explicitly model incentive constraints but instead think of satisfying any relevant incentive constraints for activity i as part of its cost c_i. For instance, this could include effort costs or Myersonian information rents.

13 A simple way to incorporate diminishing returns into the model is to have activities come in multiple units with decreasing v and/or increasing r over units.

3.1 Formalizing the $R \leq 1$ approach

3.1.1 Pre-Virus Utilitarian Objective and Definition of R_0.

Pre-virus, society solves the program

$$\max_{x \in X} \sum_{i=1}^{n} x_i v_i \tag{1}$$

Since by construction each activity has positive socioeconomic value pre-virus, society fully engages in all activities. Define

$$V_{pre\text{-}virus} \equiv \sum_{i=1}^{n} v_i$$

as the social welfare level in pre-virus society. We can define R_0

$$R_0 \equiv \sum_{i=1}^{n} r_i$$

as the virus reproduction rate in a society that engages in all of the same activities as it would pre-virus. That is, R_0 represents the reproduction rate in a society that is both fully open and that does not take even the simplest virus precautions.

3.1.2 Pure Medical Objective ("Minimize the Virus").

A society whose only goal is to minimize the spread of the virus solves the program

$$\min_{x \in X} \sum_{i=1}^{n} x_i r_i \tag{2}$$

Society thus engages only in activities with zero disease-transmission risk, i.e., with $r_i = 0$. Program (2) could be augmented to capture the idea that there is a minimal required set of essential activities, by adding a constraint $x \geq \underline{x}$, where $\underline{x}$ denotes this societal minimum set of activities.

3.1.3 Maximize Social Welfare subject to $R \leq 1$ (Proposed Paradigm).

This paper proposes the paradigm:

$$\max_{x \in X} \sum_{i=1}^{n} x_i v_i$$

subject to

$$\sum_{i=1}^{n} x_i r_i \leq 1 \tag{3}$$

The objective in program (3) is the same economic objective as in the pre-virus economy in program (1), but the constraint $\sum_{i=1}^{n} x_i r_i \leq 1$ encodes the health objective that the virus is contained. As shown above in Fig. 1, a society that imposes $R \leq 1$ as a constraint approximates the pure medical objective in (2).

3.2 Maximize social welfare subject to $R \leq 1$: solution

Let

$$\rho_i = \frac{v_i}{r_i} \tag{4}$$

denote the ratio of socioeconomic value (i.e., socioeconomic benefits minus costs) to disease-transmission risk for each activity i. That is, ρ_i represents activity i's socioeconomic value per unit of disease-transmission risk. For activities with $r_i = 0$ define $\rho_i = \infty$.

The optimal solution to (3) is found by choosing activities in descending order of their ρ_i ratios until the disease-transmission constraint is reached. Intuitively, $R \leq 1$ is a "disease-transmission budget constraint", and r_i is the "risk price" of activity i. The way to maximize social and economic well-being subject to the transmission budget constraint is to choose the activities with the highest v_i per unit of virus risk r_i, i.e., the activities with the highest ratios ρ_i. This is the standard logic of the divisible-goods version of the knapsack problem in operations research. Formally:

Proposition 1 *Let ρ_i^* denote the threshold at which the budget constraint is reached when choosing activities in descending order of ρ_i as defined in equation (4), i.e., the solution to*

$$\rho^* = \inf \left\{ \rho_i : \sum_{j:\rho_j > \rho_i} r_j \leq 1 \right\}. \tag{5}$$

The optimal choice of the activity vector x^* is:

$$x_i^* = \begin{cases} 1 & \text{if } \rho_i > \rho^* \\ q & \text{if } \rho_i = \rho^* \\ 0 & \text{if } \rho_i < \rho^*, \end{cases}$$

where $q := \left(1 - \sum_{j:\rho_j > \rho^*} r_j\right) / \left(\sum_{j:\rho_j = \rho^*} r_j\right)$ is a constant defined to exhaust the risk budget given that all activities with $\rho_i > \rho^*$ are done in full and all activities with $\rho_i < \rho^*$ are fully dropped.

The proof that the greedy solution is optimal for the fractional knapsack problem is standard and omitted (see, for example, Cormen et al. 2022, pgs. 567–568).

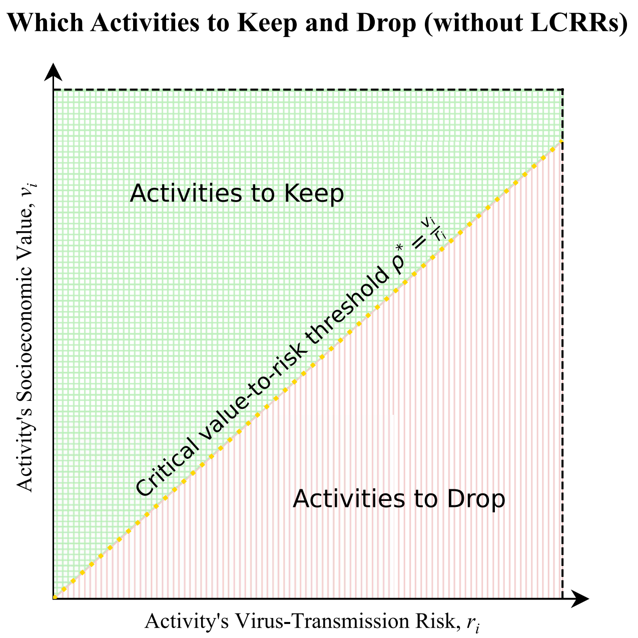

Fig. 3 *Note:* This figure illustrates the optimal mix of activities to maximize social welfare subject to $R \leq 1$. The diagonal line depicts the critical threshold ρ^* for the ratio of value-to-risk. Activities with value-to-risk ratios above ρ^* should optimally be kept and activities with value-to-risk ratios below ρ^* should optimally be dropped. The placement of the line is based on the numerical example in Sect. 5 with $R_0 = 2.5$, no LCRRs, and uniform distributions of value and risk

3.3 Maximize social welfare subject to $R \leq 1$: graphical depiction

Figure 3 presents a simple graphical depiction of the optimal solution to (3) in terms of which activities the social planner keeps and drops.

The vertical axis represents the socioeconomic value of an activity, $v_i = b_i - c_i$. The horizontal axis represents the virus-transmission risk of an activity, r_i. The diagonal line represents the critical threshold ρ^* for the ratio of value-to-risk. All activities above the diagonal line are kept (hatched green), and all activities below the diagonal line are dropped (striped red). Intuitively, activities above the diagonal line have enough "bang for buck"—enough social value per unit of virus-transmission risk—to be included in the optimum. Activities below the diagonal line are too expensive in virus-transmission terms to be worth it.

Notice that activities with very high risk can be included in the optimal solution if their socioeconomic value is sufficiently high. And, conversely, some activities with relatively low risk should be dropped if their socioeconomic value is sufficiently low. The key for the optimum is to sort activities not by their absolute level of risk, but by their ratio of value to risk.

3.4 Why treating $R \leq 1$ as a constraint is approximately optimal

This subsection proves formally why treating $R \leq 1$ as a constraint is approximately optimal.

Introduce a health cost function $h : [0, R_0] \rightarrow \mathbb{R}^+$ that takes as input a reproduction rate and outputs society's health cost. For example, the health cost could be interpreted as the value of a statistical life (VSL) times the number of deaths in a 12-month period. Figure 1 shows that this health cost is very low over the region from $R = 0.0$ to $R = 1.0$, sharply increases starting at around $R \approx 1.0$, and is then increasing and concave from $R \approx 1.0$ to $R = R_0$.

Next, consider the social planner problem

$$\max_{x \in X} \sum_i x_i v_i - h\left(\sum_i x_i r_i\right) \tag{6}$$

This program adds health costs directly to the usual socioeconomic objective function in (1), which is more standard than my proposed approach of adding $R \leq 1$ as a constraint in program (3). Our goal in this subsection is to give conditions under which the solution to (3) is an approximately optimal solution to (6).

3.4.1 Defining the benefit and cost of risk reduction curves

It will be convenient to work with benefit and cost of risk reduction curves as depicted in Fig. 2. Let $\Delta \in [0, R_0]$ denote an amount of risk reduction. We can define the benefits of risk reduction function by $b(\Delta) = h(R_0) - h(R_0 - \Delta)$.

To define the cost of risk reduction curve, let V_K be the value in the optimal solution to program:

$$\max_{x \in X} \sum_{i=1}^{n} x_i v_i$$
$$\text{subject to}$$
$$\sum_{i=1}^{n} x_i r_i \leq K$$

for any $K \in [0, R_0]$. This is the same program as (3) but with a disease-transmission budget of K instead of 1. We can define the cost of risk reduction function by $c(\Delta) = V_{pre\text{-}virus} - V_{(R_0 - \Delta)}$.

3.4.2 Convex cost of risk reduction

It is straightforward to show formally that:

Proposition 2 *The cost of risk reduction function $c(\Delta)$ is increasing and convex in the quantity of risk reduction Δ.*

Intuitively, one has to drop more and more attractive activities the larger is the desired reduction in transmission, with attractiveness defined in terms of the value-to-risk ratio $\rho_i = \frac{v_i}{r_i}$. Please see Appendix A.1 for a proof.

3.4.3 Kinked benefits of risk reduction function

The $h(\cdot)$ function shown in Fig. 1 and its corresponding $b(\cdot)$ function shown in Fig. 2 are derived from the SIR model with an initial infection seed I_0 of from 1,000 to 1,000,000 and a 12-month time horizon. It will be mathematically convenient to work with a limiting case in which the infection seed I_0 grows small and the time horizon grows long. This limiting case has an exact kink at $R = 1$.

Formally, fix a population size POP, infection fatality rate IFR, and value of statistical life VSL. Define the limiting case for the $b(\cdot)$ function as follows:

Definition 1 Let function $\bar{\pi} : [0, R_0] \to \mathbb{R}^+$ be computed as the proportion of the population infected in the SIR model in the limiting case of $I_0 \to 0$ and an infinite horizon. Define function $\bar{h} : [0, R_0] \to \mathbb{R}^+$ by $\bar{h} = POP \times IFR \times VSL \times \bar{\pi}$, the dollar value of lives lost in this limit. Define the *kinked benefits of risk reduction function* $\bar{b} : [0, R_0] \to \mathbb{R}^+$ according to $\bar{b} = \bar{h}(R_0) - \bar{h}(R_0 - \Delta)$.

Remark 5 Function $\bar{b}(\cdot)$ has the properties

(i) $\bar{b}(0) = 0$.
(ii) $\bar{b}(\Delta)$ is convex and increasing on the range $[0, R_0 - 1]$.
(iii) $\bar{b}(\Delta)$ is constant on the range $[R_0 - 1, R_0]$.

For our approximate optimality result below we need to define what it means for the true $b(\cdot)$ function to approximate the stylized $\bar{b}(\cdot)$ function with the exact kink at $R = 1$.

Definition 2 Function $b(\cdot)$ is said to ϵ*-approximate* the kinked benefits of risk reduction function $\bar{b}(\cdot)$ if $|b(\Delta) - \bar{b}(\Delta)| < \epsilon$ for all $\Delta \in [0, R_0]$.

3.4.4 Formal approximate optimality result

We are now ready to state our formal approximate optimality result.

Proposition 3 *Assume that the cost of risk reduction function $c(\cdot)$ satisfies either of the following two conditions relative to the kinked benefits of risk reduction function $\bar{b}(\cdot)$:*

(i) The marginal cost of mitigation at $R \leq 1$ (i.e., at $\Delta = R_0 - 1$) is bounded by: $c'(R_0 - 1) < \frac{\bar{b}(R_0) - \bar{b}(0)}{R_0 - 1}$. In words, the marginal cost of mitigation at $R \leq 1$ is lower than the average value of the health benefits of completely avoiding the pandemic relative to the worst case of completely ignoring the virus. ["Mitigation is sufficiently valuable"]

Or

(ii) Both (a) it is weakly optimal to do at least some mitigation: $c'(0) \leq \bar{b}'(0)$ ["Ignore is suboptimal"]; and (b) the acceleration of the cost of risk reduction curve

is smaller than the acceleration of the benefits of risk reduction curve on the interval $[0, R_0 - 1]$*, that is,* $c''(x) \leq \bar{b}''(x)$ *for* $x \in [0, R_0 - 1]$*. ["Mitigation costs do not accelerate too fast"]*

Then: a solution to (6) with $\sum_i x_i^* r_i = 1$ *is exactly optimal for the kinked benefits of risk reduction function* $\bar{b}(\cdot)$*. If the true benefits of risk reduction function* $b(\cdot)$ ϵ*-approximates the kinked benefits of risk reduction* $\bar{b}(\cdot)$*, a solution to (6) with* $\sum_i x_i^* r_i = 1$ *is approximately optimal to within* 2ϵ *of the optimal solution.*

Please see Appendix A.2 for a proof.

In words: Proposition 3 tells us that imposing $R \leq 1$ as a constraint yields a solution that approximates the optimum if either of two conditions hold. First, the marginal cost of mitigation at $R \leq 1$ is small in relation to the health benefits of completely eliminating the virus versus the worst case of completely ignoring the virus (the "mitigation is sufficiently valuable" condition). Or, second, that it is locally sub-optimal to ignore the virus and the acceleration of the cost of mitigation curve is smaller than the acceleration of the benefits of risk reduction curve up to $R \leq 1$ ("ignore is suboptimal" and "mitigation costs do not accelerate too fast").

I emphasize that in the numerical example considered below, both conditions obtain by a decent margin even without LCRRs.[14] Since LCRRs lower the marginal cost of risk reduction the conditions hold even more strongly with their use.

4 Low-cost risk reducers (LCRRs)

As I emphasized in the April 1st 2020 draft of this paper: (i) R_0 as measured, of around 2.5, describes the growth of Covid-19 in a population that is completely unaware of the virus, so is not taking even the most basic precautions against its spread; and (ii) we knew a lot about how Covid-19 spreads even very early on (e.g., see Gawande (2020) for a contemporaneous account from March 2020). Common sense thus suggested at the time, and was subsequently borne out by a variety of evidence, that some relatively simple and cheap targeted risk reductions (again, with "simple" and "cheap" understood in relation to the costs of lockdown or ignore), combined with targeted activity bans for activities with particularly poor $\frac{v}{r}$ ratios, could have generated the required 60% reduction in transmission from $R = 2.5$ to $R = 1$ without a severe societal lockdown. We did not need to eliminate all risk. We just needed to engineer a 60% reduction of risk (i.e., $\frac{2.5-1}{2.5} = 0.6$). Paul Romer (Romer 2020a, b, c) and John Cochrane (Cochrane 2020a) made this point loud and clear as well early in the pandemic.

This section models low-cost risk reducers (LCRRs) formally and describes their optimal implementation.

[14] Conditions (i) and (ii) can be confirmed visually by reference to Fig. 7, comparing the curve labeled "Benefits of Risk Reduction" to the curve labeled "Cost of Risk Reduction: No LCRRs." For condition (i), the slope of $\frac{\bar{b}(R_0)-\bar{b}(0)}{R_0-1}$ exceeds the slope of $c'(R_0 - 1)$ by about 22%. For condition (ii) the cost curve has an exponent of exactly 2.0 (driven by the uniform distribution of values and costs) and the benefits curve is approximated by a function with an exponent of approximately 2.31 on the range $\Delta \in [0, R_0 - 1]$.

4.1 Adding LCRRs to the model

We can incorporate LCRRs into the model as follows. For each activity i with $r_i > 0$, there is an LCRR technology that:

- Reduces the net socioeconomic value of the activity from v_i to $\hat{v}_i \leq v_i$
- Reduces the disease-transmission risk of the activity from r_i to $\hat{r}_i \leq r_i$

The terminology LCRRs is meant to represent any kind of intervention that reduces the socioeconomic value of an activity in exchange for reducing its disease-transmission risk. For example, facemasks are uncomfortable to wear and reduce the quality of social interactions, but can significantly reduce the risk of transmitting Covid-19 in virus in crowded indoor environments.[15] Rapid tests take time and cost money, but can significantly reduce the risk of someone who is unaware that they are infected inadvertently exposing others (Taipale et al. 2020).[16] Six feet of social distance makes factories less efficient and social life less joyful. Open windows can make rooms cold.

The reason I use the term "Low-Cost Risk Reducers (LCRRs)" rather than "Non-Pharmaceutical Interventions" (NPIs) is that the term NPIs has come to include both the low-cost, simple interventions of the sort I have in mind here as well as severe lockdowns (Ferguson et al. 2020).

4.2 Optimal LCRRs

In this section I provide simple necessary and sufficient conditions for it to be optimal to adopt LCRRs for a given activity, and then provide a formula that describes the optimal such interventions.

Proposition 4 *Let $\hat{\rho}_i = \frac{\hat{v}_i}{\hat{r}_i}$ denote activity i's value-to-risk ratio with LCRRs, analogous to ρ_i as defined in (4) without LCRRs. A necessary condition for it to be optimal to use the LCRR version of activity i is that LCRRs improve the activity's value-to-risk ratio:*

$$\hat{\rho}_i \geq \rho_i \tag{7}$$

Intuitively, LCRRs must increase society's "bang per buck" per unit of virus risk, allowing the social risk budget to be stretched further, for them to be a good idea. For

[15] In laboratory environments, N95 masks are estimated to block as much as 99.98% of viral droplets, with 95% reduction possible for cloth masks (Ma et al. (2020); see also Konda et al. (2020) for related evidence). Chu et al. (2020) provide a detailed meta-study of the medical literature on masks. In health care settings, masks are estimated to have a "relative risk" of 0.30, which corresponds to a 70% reduction in R. In non health care settings their estimate is a relative risk of 0.56 which corresponds to a 44% reduction in R. See also Abaluck et al. (2020) and Howard et al. (2021) and a summary of the available evidence at the time in Section 4.3 of Budish (2020b).

[16] Paul Romer's March 2020 blog posts showed how large-scale random testing, with isolation based on test results, can achieve the same reduction in disease transmission as a society wide lockdown at much lower costs (Romer 2020a, b, c). Under optimistic assumptions, population-scale testing on its own can get society to $R \leq 1$. See especially Figure 3 of Taipale et al. (2020). See also Droste et al. (2020) who emphasize the importance of adherence.

example, high-quality facemasks in crowded indoor settings reduced risk significantly, so that LCRR would likely pass the necessary condition. In contrast, facemasks in low-density outdoor environments (such as parks) likely did not reduce risk at all, so that would fail the necessary condition. See Appendix A.3.2 for a formal statement of the necessary condition and a proof.

Interestingly, this condition is not quite sufficient. To see why, consider the following simple example. There are two activities. Activity 1 has parameters $v_1 = 1, r_1 = 1$ without LCRRs and $\hat{v}_1 = 0.7, \hat{r}_1 = 0.5$ with LCRRs. Activity 2 has parameters $v_2 = 0.1, r_2 = 1$ without LCRRs and $\hat{v}_2 = 0.1, \hat{r}_2 = 0.5$ with. The optimum without LCRRs is to do just activity 1 (exactly exhausting the risk budget) whereas the optimum with LCRRs is to do both activities in full. Yet, utility is higher without LCRRs than with. What drives this example is that the harm of the LCRR to the more socially valuable activity (Activity 1) is relatively large, and the risk budget that is freed up is then spent on a much lower social value activity (Activity 2). Thus, in this example, even though LCRRs allow for more activity in total, welfare is lower.

The example suggests that for LCRRs to increase welfare, the marginal activities that are enabled by LCRR adoption must be of sufficiently high value relative to the utility harm of LCRRs. This intuition can be formalized as follows:

Proposition 5 *Define $\Delta r_i = r_i - \hat{r}_i$ and $\Delta v_i = v_i - \hat{v}_i$. A sufficient condition for it to be optimal to use the LCRR version of activity i is that the necessary condition (7) holds, and, additionally:*

$$\underline{\rho}_i^* \geq \frac{\Delta v_i}{\Delta r_i} \tag{8}$$

where $\underline{\rho}_i^*$ denotes the value-to-risk ratio of the marginal activity if an LCRR is adopted for activity i, taken as a lower bound over potential LCRR policies for activities other than i.

The right-hand-side of (8) represents the cost of freeing up additional risk budget by adopting LCRRs for activity i: the numerator is the utility harm of the LCRR, the denominator is the amount of risk budget that is freed up. The condition thus requires that the marginal use of the freed-up risk budget is guaranteed to be high enough to justify the utility cost of the LCRR.[17] See Appendix A.3.3 for a formal statement and proof, as well as a version of the sufficient condition that applies to a set of activities.

Now suppose that we can flexibly design the set of LCRRs for activity i. What is the optimal such set of interventions?

Proposition 6 *Assume that society's marginal value-to-risk ratio ρ^* is exogenous to the LCRR policy of activity i. Let there be K_i potential LCRR policies for activity i*

[17] To see the relationship between the necessary condition (7) and the sufficient condition (8), rearrange (7) as $\hat{\rho}_i \geq \frac{\Delta v_i}{\Delta r_i}$. (This takes several algebraic manipulations, see Appendix A.3.5). The difference is $\hat{\rho}_i$ on the left-hand-side in this manipulated version of the necessary condition, in place of $\underline{\rho}_i^*$ on the left-hand-side of the sufficient condition. Intuitively, while the sufficient condition requires that the *marginal* use of risk budget is higher than the cost of freeing up this risk budget, the necessary condition requires that an inframarginal use of risk budget is higher. Since this manipulated version of the necessary condition is not very interpretable, I prefer to use $\hat{\rho}_i \geq \rho_i$ as the necessary condition.

with value and risk $(v_i^k, r_i^k)_{k=1}^{K_i}$. *Let* $\Delta v_i^k = v_i - v_i^k$ *and* $\Delta r_i^k = r_i - r_i^k$. *The optimal LCRR policy for activity* i *is the choice that maximizes:*

$$\underbrace{\Delta r_i^k}_{\substack{\textit{risk reduction}\\ \textit{from LCRR } k}} \cdot \underbrace{\rho^*}_{\substack{\textit{marginal value}\\ \textit{of risk budget}}} - \underbrace{\Delta v_i^k}_{\substack{\textit{utility harm}\\ \textit{of LCRR } k}} \tag{9}$$

In words: the optimal LCRR policy for activity i maximizes its risk reduction (Δr_i) times the marginal societal value of this additional risk budget (ρ^*) minus the utility harm of the LCRRs (Δv_i). While intuitive, notice that the optimal LCRR policy for activity i is not necessarily the one that maximizes the ratio of value-to-risk. Eliminating the last epsilon of risk is great for the ratio of value-to-risk but only has a small benefit for social welfare. This may not be worth it if the utility cost of eliminating this last epsilon of risk is high. Instead, (9) tells us that the optimal LCRR policies are those that achieve large absolute reductions of the quantity of risk, at small absolute harm to utility.

4.3 Graphical intuition

Figure 4 illustrates the price-theory intuition for the effect of LCRRs on optimal policy. Under the conditions described above, LCRRs increase the social welfare that is achievable for any given level of disease-transmission risk. This is equivalent to LCRRs reducing the economic cost of virus mitigation. The figure illustrates a case where, without LCRRs, mitigation to $R \leq 1$ is optimal but expensive. With LCRRs, $R \leq 1$ remains optimal but is significantly less expensive. Specific numerical examples of mitigation cost curves, and how they are impacted by LCRRs, will be provided in Sect. 5.

The figure may represent a society that is initially in lockdown, at significant economic expense, and then transitions to reopening with LCRRs such as rapid tests and facemasks that can help keep $R \leq 1$ at much lower cost until a vaccine is available.

5 Detailed numerical example: the enormous social value of low-cost risk reducers

This section presents a detailed numerical example in two steps. First, it considers the baseline model of Sect. 3 and depicts the optimal activity mix and cost of mitigation curve, in a simple numerical environment using uniform distributions and empirical evidence on R_0. Second, I add LCRRs to the example. This yields the most important results of the section: magnitudes for just how economically valuable LCRRs are, given what we know about their efficacy and R_0.

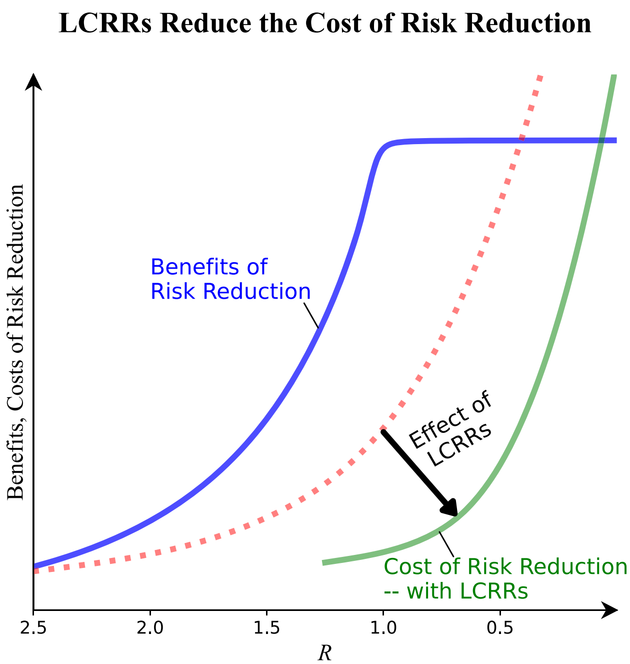

Fig. 4 *Note:* The solid-blue line and dotted-red line are the same benefits and costs of risk reduction curves as in Fig. 2. The solid-green line illustrates how low-cost risk reducers (LCRRs) lower the economic cost of risk reduction. The illustrated reduction is based on the numerical example from Sect. 5 under the assumption that LCRRs reduce risk by 50%

5.1 Numerical example: initial model without LCRRs

Let activities' values and risks be jointly uniformly distributed on $[0, \bar{v}] \times [0, \bar{r}]$. If society engages in all activities, i.e., society engages in Program (1) above, then pre-virus social welfare is given by

$$V_0 = \int_0^{\bar{r}} \int_0^{\bar{v}} v_i \frac{1}{\bar{v}\bar{r}} dv dr = \frac{\bar{v}}{2},$$

and the reproduction rate of the virus without any LCRRs or activity restrictions is given by

$$R_0 = \int_0^{\bar{r}} \int_0^{\bar{v}} r_i \frac{1}{\bar{v}\bar{r}} dv dr = \frac{\bar{r}}{2}.$$

Without loss of generality, set $\bar{v} = 1$. The relation $\bar{r} = 2R_0$ allows us to express $\bar{r}$ based on a parameter choice for R_0, which can be based on empirical evidence.

Now consider the constrained problem in which society chooses activities to maximize social welfare subject to a constraint on R; particular attention will be given to the constraint $R \leq 1$. Formally, for any $K \in [0, R_0]$, society solves

Table 2 Optimal Solution to Achieve $R \leq 1$, without LCRRs

	Value of R_0				
	2.0	2.5	3.0	3.5	4.0
To Achieve $R \leq 1$:					
% Activities Dropped	37.5	45.0	50.0	53.7	56.7
% Social Welfare Dropped	18.8	27.0	33.3	38.3	42.3
Relative to Pre-Virus Economy:					
% Activities Kept	62.5	55.0	50.0	46.3	43.3
% Social Welfare Kept	81.2	73.0	66.7	61.7	57.7

Please see the text of Sect. 5.1 for description of the numerical example without LCRRs

$$\max_{x(\cdot)} \int_0^{2R_0} \int_0^1 x(v,r) \cdot v_i \frac{1}{2R_0} dvdr$$
$$\text{subject to}$$
$$\int_0^{2R_0} \int_0^1 x(v,r) \cdot r_i \frac{1}{2R_0} dvdr \leq K \tag{10}$$

The analysis in Sect. 3 shows that the optimal solution to (10) is characterized by a value-to-risk threshold ρ^*, such that all activities with value-to-risk ratio above ρ^* are included and all activities with value-to-risk ratio below ρ^* are dropped. Appendix A.4 obtains ρ^* for this example in closed form.

Table 2 describes the features of the optimal solution to achieve $R \leq 1$, for values of R_0 ranging from 2.0 to 4.0. Let me highlight the results with $R_0 = 2.5$, which was the CDC's best estimate as of Sept 2020 and was the midpoint of the Imperial study's range in March 2020. In this case, to achieve $R \leq 1$ requires dropping 45% of activities that together constitute 60% of risk and 27% of social welfare. Society keeps the other 55% of activities, leaving it with just 73% of pre-virus social welfare. Even though society keeps the activities with the highest value-to-risk ratio, the welfare cost is significant: over one-quarter of social welfare.

Even at the CDC's optimistic scenario, $R_0 = 2.0$, achieving $R \leq 1$ requires dropping 38% of activities constituting 19% of social welfare. At the CDC's pessimistic scenario, $R_0 = 4.0$, achieving $R \leq 1$ requires dropping 57% of activities constituting 42% of social welfare.

Remark: Super-Spreader Activities.

A limitation of the uniform-distribution assumption is that it does not allow for super-spreader activities—a small mass of activities with particularly large r_i. Intuitively, if super-spreader activities are incorporated into the example, then significantly fewer activities need to be dropped to get to $R \leq 1$, so social welfare can be significantly higher than analyzed here (unless the super-spreader activities are also an unusually large fraction of pre-virus social welfare).

Table 3 Optimal Solution to Achieve $R \leq 1$: Large Effect of LCRRs

	LCRR Efficacy					
	No LCRRs	30%	40%	50%	60%	70%
To Achieve $R \leq 1$:						
% Activities Dropped	45.0	32.1	25.0	15.0	0.0	0.0
% Social Welfare Dropped	27.0	13.8	8.3	3.0	0.0	0.0
Relative to Pre-Virus Economy:						
% Activities Kept	55.0	67.9	75.0	85.0	100.0	100.0
% Social Welfare Kept	73.0	86.2	91.7	97.0	100.0	100.0

Please see the text of Sect. 5.1-5.2 for description of the numerical example with LCRRs. The $R_0 = 2.5$ scenario is based on the CDC's current best estimate

5.2 LCRRs

Now add LCRRs to the example. For simplicity, assume that LCRRs reduce risk by a uniform percentage across activities, denoted γ_r, and similarly reduce socioeconomic value by a uniform percentage across activities, denoted γ_v. Thus activity i with original value v_i and risk r_i has with-LCRR value of $\hat{v}_i = (1 - \gamma_v)v_i$ and risk of $\hat{r}_i = (1 - \gamma_r)r_i$. With this assumption we still have a joint uniform distribution of value and risk, only now on $[0, (1 - \gamma_v)\bar{v}] \times [0, (1 - \gamma_r)\bar{r}]$. Thus all of the same math from above goes through analogously.[18]

Table 3 summarizes the analysis. The first column repeats the figures without LCRRs focusing on $R_0 = 2.5$ as the baseline case. The remaining columns vary LCRR efficacy from 30%-70%. The studies cited in Sect. 4 suggest that 50% is a reasonable ballpark estimate for the use of either high-quality facemasks or widescale rapid testing. I use 30% as a conservative figure for either LCRR technology, and I use 70% to represent the effective use of the full suite of LCRR technologies.

Focus first on the 50% efficacy column. LCRRs alone reduce average transmission from $R_0 = 2.5$ to $R = 1.25$ without any reduction in activity levels. Thus, to achieve $R \leq 1$ requires a further 20% reduction of risk. This can be accomplished by dropping the 15% of activities with the worst value-to-risk ratios, which together constitute just 3% of pre-virus social welfare. Society maintains activities that together constitute 97% of pre-virus social welfare, gross of the cost of the LCRRs.

If LCRRs are sufficiently effective, they alone can reduce R to less than 1 without any reduction in activity levels. This occurs if $R_0(1-\gamma_r) \leq 1$; for example, if $R_0 = 2.5$ and the reduction in risk is at least 60%. If we think beyond the uniform-distribution example to include a small mass of super-spreader activities, as discussed earlier, then the more likely practical conclusion is that society would only have to drop super-spreader activities, keeping everything else.

In a more pessimistic scenario with $R_0 = 4.0$, LCRRs can have an especially dramatic effect on social welfare. Without LCRRs, society has to drop over 50% of

[18] A nuance is that for activities with very high value and very high risk it might be optimal not to use the LCRR, as discussed in Sect. 4.2. The figures in Table 3 can thus be interpreted as a lower bound on the social welfare benefits of LCRRs.

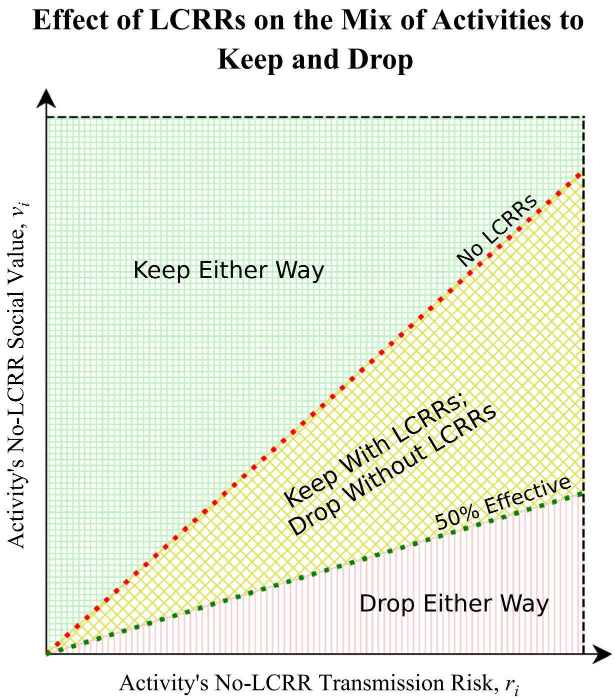

Fig. 5 *Note:* This figure illustrates how LCRRs expand the optimal set of activities used to maximize societal welfare subject to $R \leq 1$. The diagonal line labeled "No LCRRs" depicts the critical threshold ρ^* for the ratio of value-to-risk if there are no LCRRs; this is the same line as in Fig. 3 and is based on the uniform-distribution example with $R_0 = 2.5$. Without LCRRs, activities above the "No LCRRs" line should optimally be kept and activities below this line should optimally be dropped. The diagonal line labeled "50% Effective" depicts how the mix of activities expands if LCRRs are adopted and they uniformly reduce risk by 50%. Activities with social value and risk in the yellow hatched region, above the "50% Effective" line but below the "No LCRRs" line, can be included in the optimum with LCRRs whereas they are dropped without

activity and over 40% of social welfare to reach $R \leq 1$. If LCRRs are 60% effective, society can drop 28% of activity constituting 10% of social value. If LCRRs are 70% effective, society can drop just 12.5% of activity constituting just 2% of pre-virus social value.

Figure 5 illustrates how LCRRs affect the optimal activity mix. The top-left region (green squares) depicts activities that are included in the optimum whether or not LCRRs are utilized. If LCRRs are utilized, these activities are lower risk, and society then optimally spends this freed-up risk budget on the yellow-hatched region labeled "Keep With LCRRs, Drop Without LCRRs". These are the activities that, without LCRRs, are too risky per unit of utility, but with LCRRs can be included in the optimum. The bottom right red-striped region depicts activities that are optimally dropped even with LCRRs; these are the activities with the lowest value-to-risk ratios. The figure depicts the dividing line for the case of $R_0 = 2.5$ and 50% LCRR effectiveness. The higher is LCRR effectiveness, the smaller is this striped region, vanishing to zero if the risk reduction is 60% or greater.

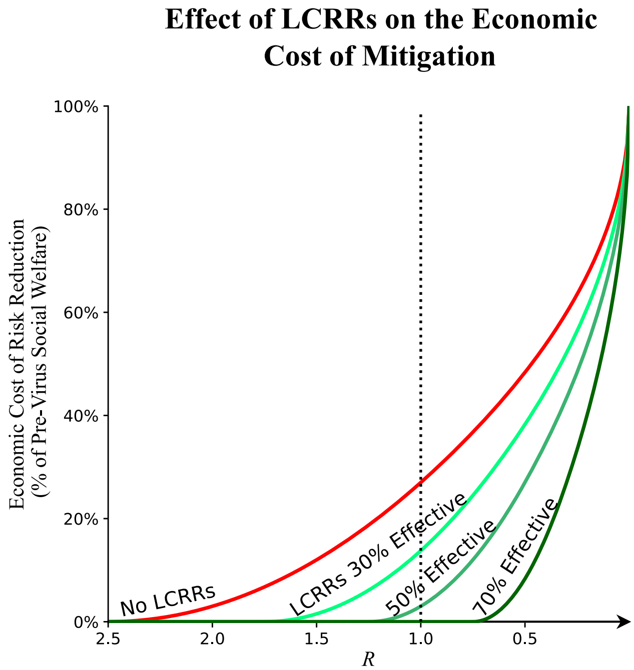

Fig. 6 *Note:* This figure presents the economic cost of risk reduction curve in the numerical example with no LCRRs and with LCRRs of varying effectiveness. The horizontal axis represents the level of mitigation and the vertical axis is the cost depicted as a percentage of pre-virus social value $V_{pre-virus}$ as defined above. These curves correspond to the illustrative cost of risk reduction curves presented in price-theory diagrams Fig. 2 (without LCRRs) and Fig. 4 (with LCRRs)

Figure 6 illustrates the effect of LCRRs on the cost of mitigation curves. This figure complements the illustrative price theory diagram provided earlier, Fig. 4. Focus first on the 50% effectiveness line. LCRRs get us from $R_0 = 2.5$ to $R = 1.25$ without dropping any activities. Then, the cost curve begins increasing as activities are dropped, but at first this increase is slow because society drops only those activities with the poorest value-to-risk profiles. For this reason, society can get all the way to $R = 1$ very cheaply. If LCRRs are 30% effective, the cost is meaningfully higher but still much lower than without LCRRs. If LCRRs are 70% effective, society can get all the way to $R = 1$ for free. Yet, even in this optimistic case, a policy of "minimize the virus" remains very expensive—the costs get arbitrarily high as society drops more and more activities with non-zero risk.

6 Dynamic considerations

In the model considered so far, $v_i = b_i - c_i$ represents the socioeconomic value of an activity ignoring any virus considerations. The virus enters the analysis through r_i, the risk of transmission; the shadow cost of the $R \leq 1$ constraint is a way of capturing the negative externalities of otherwise socially valuable behavior that risks spreading the virus.

Clearly, if the level of infections in an area is high enough, this will more directly affect the benefits and costs of many kinds of activities because of the fear of catching the virus. Goolsbee and Syverson (2021) provide empirical evidence on the quantitative importance of this channel.

This idea can be captured in the model by letting $v_i^{exposed}$ denote the utility from activity i if perceived exposure to the virus is high, and assuming that $v_i > v_i^{exposed}$ for all i. If $v_i^{exposed} < 0$ for some activity then individuals will drop the activity on their own, even without any kind of formal ban, as documented by Goolsbee and Syverson (2021). Since the high-exposure value $v_i^{exposed}$ is worse than the original value v_i for all activities (whether positive or negative), the level of feasible social welfare is lower for any target constraint R. Formally, for any $K \in [0, R_0]$, define V_K as in Sect. 3.2 and define $V_K^{exposed}$ as the value in the optimal solution to program:

$$\max_{x \in X} \sum_{i=1}^{n} x_i v_i^{exposed}$$

subject to

$$\sum_{i=1}^{n} x_i r_i \leq K$$

Then we have $V_K^{exposed} < V_K$ for all $K \in [0, R_0]$. This is a simple way of articulating the value of treating $R \leq 1$ as a constraint before the stock of infections grows, on social and economic grounds alone. Several papers elaborating what are now called Behavioral SIR models make a version of this point, and several suggest that, since $v_i^{exposed}$ seems likely to decrease monotonically with the stock of infections, R may automatically equilibrate to around 1—but at a high stock of infections and with a fearful, low-utility society, what McAdams (2021) refers to as "epidemic limbo," as opposed to with a low stock of infections and higher social welfare level.[19]

A society that does not take early action and faces a high stock of infections faces a more complicated dynamic problem than a society that acts to constrain $R \leq 1$ when the stock of infections is low. In particular, a society with a high stock of infections may wish to first invest in significantly reducing the stock of infections (i.e., R significantly lower than 1), before then transitioning to a steady state with $R \leq 1$ as analyzed above. Some specific dynamic scenarios will be discussed in the next section.

[19] See Cochrane (2020b), Toxvaerd (2020), Keppo et al. (2020) and Farboodi et al. (2021) for models with this equilibration feature, and see Atkeson et al. (2024) for related stylized empirical facts.

Table 4 Sense of Magnitudes for the Potential Health and Economics Costs of the $R \leq 1$ Policy Paradigm

Initial Stock of Infections	10,000	100,000	1,000,000
Health Outcomes			
Total Infections	718,000	5,810,000	24,900,000
Total Deaths	5,000	40,000	174,000
Socioeconomic Outcomes			
% Activities Maintained	85.0%	85.0%	85.0%
% Social Welfare Maintained	97.0%	97.0%	97.0%
Overall Magnitudes			
Health Cost	\$35 B	\$282 B	\$1.22 T
Economic Cost	\$660 B	\$660 B	\$660 B
Total Cost	\$695 B	\$942 B	\$1.88 T

See the text of Sect. 7.1

7 Sense of magnitudes for the potential costs of this paper's policy relative to actual policy

7.1 Overall magnitudes

Table 4 gives a sense of magnitudes for the potential health and economic costs of the $R \leq 1$ policy paradigm advocated in this paper. The table assumes $R_0 = 2.5$ and that LCRRs can reduce risk by 50%. Columns vary the initial seed of infections I_0. The table reflects a population and economy the size of the United States, using a population of 330 million, a value of statistical life (VSL) of \$7 million, and annual GDP of \$22 trillion.

If the $R \leq 1$ policy is implemented early, cumulative deaths are very small: only 5,000 deaths in 12 months if the initial seed is 10,000 infections, and 40,000 deaths in 12 months if the initial seed is 100,000 infections. This latter figure is roughly the same as annual traffic fatalities. At a VSL of \$7 million, the health cost is about \$280 billion. Note that this health cost does not include the cost of non-fatal infections.[20] My analysis of health costs also does not attempt to account for the benefit of reducing the likelihood of variants by keeping the overall number of infections low (Yamey 2021).

The economic cost of implementing $R \leq 1$, assuming LCRRs are utilized, is that 15% of activities constituting 3% of social welfare have to be dropped in targeted activity bans. One rough way of estimating what 3% of social welfare means in dollar terms is to use 3% of GDP, which is about \$660 billion per year in the United States. Thus, the total dollar cost of the $R \leq 1$ approach if the initial stock of infections is 100,000 is around \$940 billion, not including the cost of the LCRRs.

If the $R \leq 1$ policy is implemented after the initial stock of infections has grown considerably then the number of deaths is larger. If the stock of infections is 1 million, then the cumulative number of deaths is 174,000. If the stock of infections is 2 million,

[20] Since the analysis uses an infection fatality rate of 0.7%, a simple way to incorporate the cost of non-fatal infections is to increase the VSL figure used by an amount equal to $1/.007 = 143$ times the assumed average cost of a non-fatal infection.

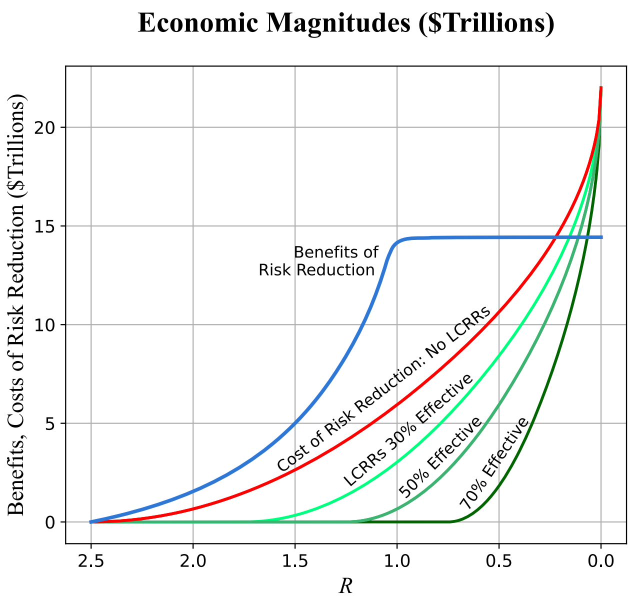

Fig. 7 *Note:* The Benefits curve is the same as in Fig. 2, with the vertical scale based on a VSL of $7 million. The Cost curves are the same as in Fig. 6 with the vertical scale set using $V_{pre-virus} = \$22$ trillion, which was US forecast GDP for 2020

which may be a rough estimate for the peak in the United States, [21] then the cumulative number of deaths is 245,000. In dollar terms, these figures correspond to about $1.2-$1.7 trillion of health costs. Thus, the total dollar cost of the $R \leq 1$ approach if implemented after the initial stock has grown large is about $1.9-$2.4 trillion.

Figure 7 presents the benefit and cost curves from earlier as another way of getting a sense of magnitudes for the optimal policy. The benefits curve uses a VSL of $7 million and the cost curve is scaled to GDP of $22 trillion. At $R = 1.0$, about $14 trillion of health benefits are achieved relative to a society that ignores the virus, at a cost of $660 billion with 50%-effective LCRRs (gross of the cost of the LCRRs).

7.2 Value of dynamic plans if the stock of infections is high

In this case of a high stock of infections there can be significant gains to a dynamic strategy, as discussed in the previous section. Table 5 considers simple dynamic plans that consist of a target for R of significantly less than 1 for a period of 30, 60 or 90 days, followed by a target of $R \leq 1$ for the remainder of the year. It focuses on the case of an initial stock of 1 million. I numerically compute the optimal R target for the early period. The longer is the period with the lower R target, the less strict this lower R target can be, and vice versa. What the early period does is invest in

[21] Daily Covid-19 deaths in the United States initially peaked at 2,752 on April 15th 2020, which using a 0.7% infection fatality rate corresponds to 393,143 new infections per day. With a 5-day infectiousness duration (i.e., $\frac{1}{\gamma} = 5$ in the SIR model discussed in Sect. 2), this is a stock of about 2 million infections.

Table 5 Potential Health and Economics Costs for Dynamic Policies

Scenario			
Duration of Initial Phase	30 days	60 days	90 days
Optimal R Target for Initial Phase	0.59	0.76	0.83
Health Outcomes			
Cumulative Infections in Phase 1	2,290,000	3,900,000	5,430,000
Stock of Infections Entering Phase 2	91,000	54,000	40,000
Cumulative Infections during Phase 2	4,160,000	2,170,000	1,370,000
Cumulative Total Infections	6,450,000	6,070,000	6,800,000
Cumulative Total Deaths	45,000	42,000	48,000
Socioeconomic Outcomes			
% Activities Kept during Phase 1	60.4%	70.6%	74.8%
% Social Welfare Kept during Phase 1	79.1%	88.5%	91.5%
Time-Weighted % Activities Kept	83.0%	82.6%	82.5%
Time-Weighted % Social Welfare Kept	95.5%	95.6%	95.7%
Overall Magnitudes			
Health Cost	\$314 B	\$297 B	\$333 B
Economic Cost	\$984 B	\$968 B	\$957 B
Total Cost	\$1.30 T	\$1.26 T	\$1.29 T

Each dynamic policy analysis assumes the initial stock of infections is 1 million. See the text of Sect. 7.2 for full details

reducing the stock of infections. In each of the optimal plans the stock is reduced from 1 million to under 100,000. This lowers the number of deaths significantly at higher cost to socioeconomic value during the period of higher restrictions. To highlight one scenario, an R target of 0.76 for 60 days reduces the stock of infections from 1 million to 54,000, so reduces cumulative total deaths to 42,000 versus the 174,000 deaths under a static $R \leq 1$ plan. This lowers health costs by about \$900 billion while raising the cost of lost social welfare by about \$300 billion, because society loses 12% of total social welfare during the severe period (30% of activities).

Thus, if the stock of infections grows to 1 million before policy action, a simple dynamic plan consisting of a short period of $R \ll 1$ followed by a longer period of $R \leq 1$ can reduce costs by about \$600 billion relative to a static plan of $R \leq 1$ throughout.[22] Additionally, the dynamic plan likely has additional economic gains in this scenario of late policy action due to the fear of the virus channel documented by Goolsbee and Syverson (2021) and discussed in Sect. 6.

7.3 Actual policy costs

Let us now compare the health and socioeconomic costs computed here to actual policy in the United States. In the 12 months from March 2020-Feb 2021 there were 512,978 deaths from Covid-19. At a \$7 million VSL, the cumulative cost of these deaths is

[22] If the stock of infections is 1,000,000 the optimal static plan is $R = 0.94$. In this plan, cumulative total deaths are 88,000, social welfare losses are 5%, and the total economic cost in 12 months is \$1.6 trillion.

about \$3.6 trillion. This corresponds to a weighted-average R of about $1.12-1.14$, reflecting some periods of significant lockdown and some periods with R significantly greater than 1. This is hundreds of thousands more deaths than an $R \leq 1$ policy under any of the scenarios considered.

In this time period there was about \$1 trillion of lost GDP and \$3.5 trillion of economic stimulus, with another \$1.9 trillion of economic stimulus passed in the 13th month (March 2021). This is trillions of dollars higher than the economic costs of an $R \leq 1$ policy under any of the scenarios in Tables 4 and 5.

There were also enormous costs to children from educational losses. Hanushek and Woessmann (2020) estimate costs on the order of \$14-\$28 trillion in the United States, and Azevedo et al. (2021) estimate costs of \$20 trillion globally, with the former estimate based on the net present value of lost GDP growth and the latter based on the net present value of lost earnings for the affected students. Spending some of society's disease-transmission budget on keeping schools open, making effective use of LCRRs to further increase schools' value-to-risk ratio, could have avoided this tragedy.

7.4 Summary

It thus seems plausible that the approach advocated in this paper could have saved hundreds of thousands of lives, trillions of dollars, and reduced severe harms to a generation of students.

8 Conclusion: a new play in the pandemic playbook

This paper has argued that the Covid-19 pandemic called for a novel play in the pandemic playbook: (i) pre-vaccine, treat $R \leq 1$ as a constraint and maximize social welfare subject to this constraint, (ii) use low-cost risk reducers and targeted activity bans to get to $R \leq 1$ as efficiently as possible, and (iii) accelerate vaccination to essentially the maximum extent possible, as studied in companion papers Ahuja et al. (2021) and Castillo et al. (2021). The gains relative to actual policy feasibly would have been in the trillions of dollars and hundreds of thousands of lives in the United States alone.

I want to conclude by contrasting the novel play called for by Covid-19 with other past plays in what we might call the "*pandemic playbook*." See Fig. 8. Panels A and B repeat Figs. 2 and 4 from earlier in the paper. In Panel A, an $R \leq 1$ policy is approximately optimal but expensive. In Panel B, the $R \leq 1$ policy is much cheaper because of the use of LCRRs.

Panel C depicts the case where it is optimal to ignore an infectious threat, because the health benefits of mitigation are not large enough to justify the costs of mitigation.

Panel D depicts the case where the optimal policy is to partially mitigate the spread to some amount R greater than 1, but where mitigation beyond this point becomes too expensive to justify further risk reduction. This could represent a scenario in which there are some interventions that are both cheap and effective—"no brainers" to do as fully as possible—but then it is impossible or prohibitively expensive to do more.

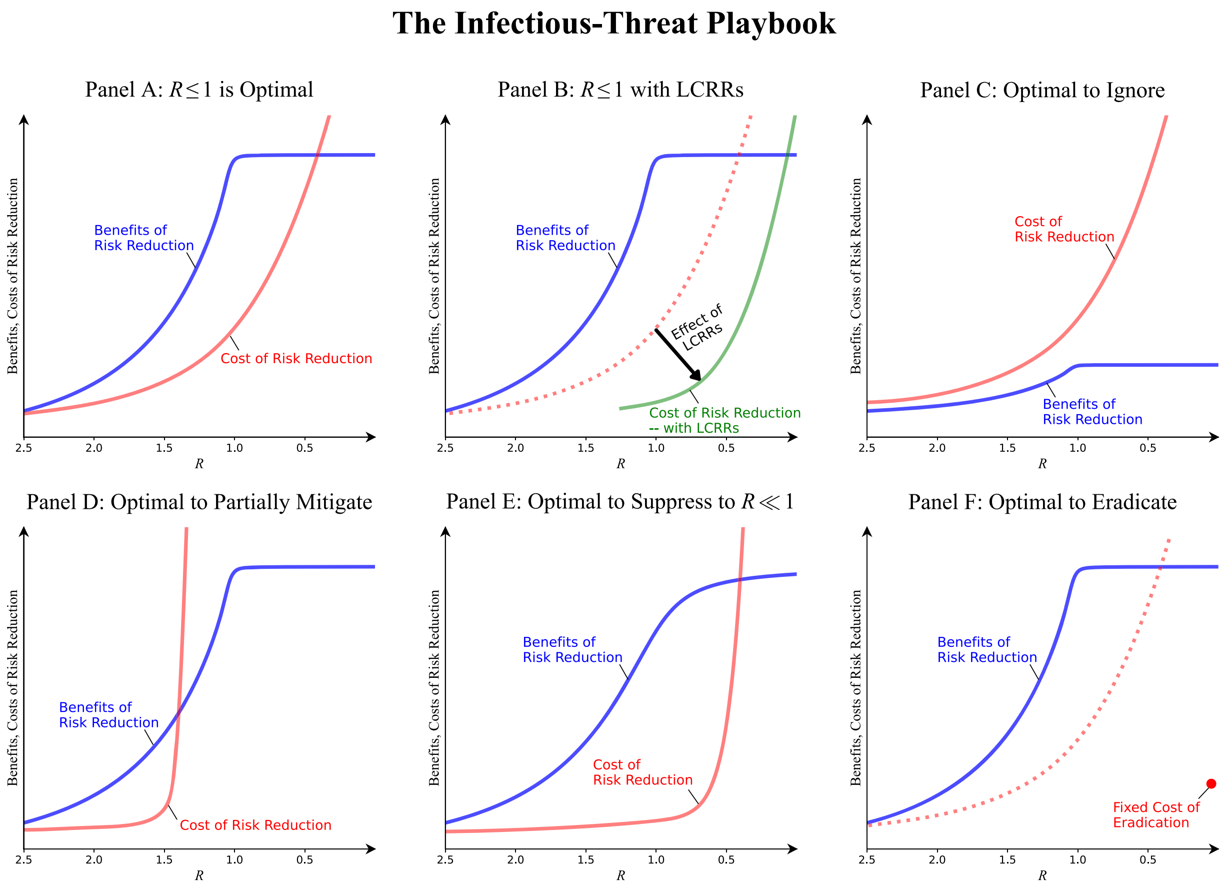

Fig. 8 *Note:* The blue line of Panels A, B, C, D, and F depicts the same information as the $I_0 = 100,000$ case of Figure 1, but with both axes flipped as described in the text of Sect. 2. The blue line of panel E uses $I_0 = 10$ million. The cost curves depicted in Panels A and B are as discussed in the body of the paper. The cost curves depicted in Panels C-F are conceptual and illustrative. Please see the text of Sect. 8 for discussion of each case

This may be a useful way to think about aspects of the AIDS pandemic. Public health officials converged on a suite of mitigation responses, including safe sex education, condom distribution, needle exchange, testing donated blood, etc. (what this paper calls LCRRs), and also, more controversially, closing bath houses in some cities (what this paper calls targeted activity bans). But these measures alone were not enough to reduce the spread to $R \leq 1$. Notably, no public health expert recommended literally trying to minimize the spread of the virus by banning sex (or banning non-monogamous sex, etc.), whereas in response to Covid-19 many public health officials were comfortable with bans on significant amounts of social contact.

Panel E depicts the case where it is optimal to suppress to R significantly less than 1 ($R \approx 0.7 - 0.8$ is optimal in the figure). This requires both (i) a large initial stock of infections (the figure uses I_0 =10 million) and (ii) that the cost of mitigation beyond $R < 1$ does not grow too rapidly. As discussed in Sects. 6-7, it can also be valuable to use a simple dynamic plan that first invests in reducing the stock of infections and then maintains $R \leq 1$ until vaccines are available.

Panel F illustrates the case of eradication. Eradication can be viewed as requiring a one-time fixed cost as opposed to an ongoing flow cost of reducing spread via LCRRs or activity bans. Eradication is optimal if this one-time fixed cost is sufficiently low. This likely was the case in some countries that successfully implemented Covid zero policies for a period of time (e.g., Australia) but likely was not the case in many countries that acted after the initial stock of infections had grown beyond the point of feasible eradication.

The epidemiologist Dr. Michael Osterholm wrote “As epidemiologists, we have two goals. The first is to prevent. When that is not possible, the second is to minimize...”(Osterholm and Olshaker 2020, pg. 26). Dr. Francis Collins, quoted in the introduction, said “If you’re a public-health person... you attach *infinite value to stopping the disease* and saving a life.”

The public-health instinct expressed by Drs. Collins and Osterholm is a useful heuristic for optimal policy in the scenarios depicted in Panels D, E, and F. Specifically, I think it is reasonable to understand the instinct to “minimize” and to “attach infinite value to stopping the disease and saving a life” to mean pursuing eradication if feasible at a fathomable fixed cost (Panel F), and to mean engaging in all of the relatively cheap interventions on the relatively flat parts of the Cost-of-Risk-Reduction curves as fully as possible, but then to stop when the curve gets much steeper (Panels D and E).

However, the public-health instinct to minimize, and to regard the economy and education as merely “collateral damage,” was profoundly suboptimal in the case depicted in Panels A and B. Covid-19 demanded a new play in the pandemic playbook.

Supplementary Information The online version contains supplementary material available at https://doi.org/10.1007/s10058-025-00379-z.

References

Abaluck J, Chevalier JA, Christakis NA, Forman HP, Kaplan EH, Ko A, Vermund SH (2020) The case for universal cloth mask adoption and policies to increase supply of medical masks for health workers. Available at SSRN 3567438

Acemoglu D, Chernozhukov V, Werning I, Whinston MD (2020) A multi-risk SIR model with optimally targeted lockdown. NBER Working Paper 27102

Ahuja A, Athey S, Baker A, Budish E, Castillo JC, Glennerster R, Kominers SD, Kremer M, Lee J, Prendergast C, Snyder CM, Tabarrok A, Tan BJ, Wiecek W (2021) Preparing for a pandemic: accelerating vaccine availability. AEA Pap Proc 111:331–35

Akbarpour M, Budish E, Dworczak P, Kominers SD (2024) An economic framework for vaccine prioritization. Q J Econ 139(1):359–417

Alvarez F, Argente D, Lippi F (2020) A simple planning problem for COVID-19 lockdown. NBER Working paper 26981

Atkeson AG, Kopecky KA, Zha T (2024) Four Stylized Facts about COVID-19. Int Econ Rev 65(1):3–42

Avery C, Bossert W, Clark A, Ellison G, Ellison SF (2020) An economist's guide to epidemiology models of infectious disease. J Econ Persp 34(4):79–104

Azevedo JP, Hasan A, Goldemberg D, Geven K, Iqbal SA (2021) Simulating the potential impacts of COVID-19 school closures on schooling and learning outcomes: a set of global estimates. World Bank Res Observ 36(1):1–40

Budish E (2012) Matching versus mechanism design. ACM SIGecom Exchanges 11(2):4–15

Budish E (2020a) $R < 1$ as an Economic Constraint: Can We 'Expand the Frontier in the Fight Against Covid-19?'. https://repec.bfi.uchicago.edu/RePEc/pdfs/BFI_WP_202031.pdf. Last revised April 1, 2020

Budish E (2020b) Maximize Utility subject to $R \leq 1$: a simple price-theory approach to covid-19 lockdown and reopening policy. NBER Working Paper No. 28093. Available at https://www.nber.org/papers/w28093. Last revised November, 2020

Castillo JC, Ahuja A, Athey S, Baker A, Budish E, Chipty T, Glennerster R, Kominers SD, Kremer M, Larson G, Lee J, Prendergast C, Snyder CM, Tabarrok A, Tan BJ, Wiecek W (2021) Market design to accelerate COVID-19 vaccine supply. Science 371(6534):1107–1109

Centers for Disease Control and Prevention (2020) COVID-19 pandemic planning scenarios. Accessed on 28 June, 2024 from https://archive.cdc.gov/#/details?q=hcp/plan&start=0&rows=10&url=https://www.cdc.gov/coronavirus/2019-ncov/hcp/planning-scenarios.html

Chu DK, Akl EA, Duda S, Solo K, Yaacoub S, Schünemann HJ, El-harakeh A, Bognanni A, Lotfi T, Loeb M et al (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 395(10242):1973–1987

Cochrane JH (2020a) Beyond testing: the central question for pandemic policy. March 29. https://johnhcochrane.blogspot.com/2020/03/beyond-testing-central-question-for.html

Cochrane JH (2020b) An SIR model with behavior. May 4. Available at https://johnhcochrane.blogspot.com/2020/05/an-sir-model-with-behavior.html

Collins F (2023) A deplorable and an elitist walk into a bar: francis collins and wilk wilkinson. Moderated by Martha Bayles, virtual panel, undated, posted July 10, 2023, by Braver Angels, YouTube, https://www.youtube.com/watch?v=W1eAvh1sWiw. Quote starting at 54:33

Cormen TH, Leiserson CE, Rivest RL, Stein C (2022) Introduction to algorithms, 4th edn. MIT Press

Droste MC, Stock J, Atkeson A (2020) Economic benefits of COVID-19 screening tests. medRxiv

Eichenbaum MS, Rebelo S, Trabandt M (2020) The macroeconomics of epidemics. NBER Working Paper

Farboodi M, Jarosch G, Shimer R (2020) Internal and external effects of social distancing in a pandemic. NBER Working paper 27059

Farboodi M, Jarosch G, Shimer R (2021) Internal and external effects of social distancing in a pandemic. J Econ Theory 196:105293

Ferguson NM, Laydon D, Nedjati-Gilani G, Imai N, Ainslie K, Baguelin M, Bhatia S, Boonyasiri A, Cucunubá Z, Cuomo-Dannenburg G, et al. (2020) Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand. Available at https://www.imperial.ac.uk/media/imperial-college/medicine/sph/ide/gida-fellowships/Imperial-College-COVID19-NPI-modelling-16-03-2020.pdf

Fernández-Villaverde J, Jones CI (2022) Estimating and simulating a SIRD model of COVID-19 for many countries, states, and cities. J Econ Dyn Control 140:104318

Flaxman S, Mishra S, Axel Gandy H, Unwin JT, Mellan TA, Coupland H, Whittaker C, Zhu H, Berah T, Eaton JW et al (2020) Estimating the effects of non-pharmaceutical interventions on COVID-19 in Europe. Nature 584(7820):257–261

Gawande A (2020) Keeping the coronavirus from infecting health-care workers. The New Yorker, March 21. Available at https://www.newyorker.com/news/news-desk/keeping-the-coronavirus-from-infecting-health-care-workers

Goldhaber D, Kane TJ, McEachin A, Morton E, Patterson T, Staiger DO (2023) The educational consequences of remote and hybrid instruction during the pandemic. Am Econ Rev Insights 5(3):377–92

Goolsbee A, Syverson C (2021) Fear, lockdown, and diversion: comparing drivers of pandemic economic decline 2020. J Public Econ 193:104311

Hanushek EA, Woessmann L (2020) The economic impacts of learning losses. OECD Education Working Papers

Howard J, Huang A, Li Z, Tufekci Z, Zdimal V, van der Westhuizen H-M, von Delft A, Price A, Fridman L, Tang L-H et al. (2021) An evidence review of face masks against COVID-19. Proc Natl Acad Sci 118(4)

Jack R, Oster E (2023) COVID-19, school closures, and outcomes. J Econ Persp 37(4):51–70

Jack R, Halloran C, Okun J, Oster E (2023) Pandemic schooling mode and student test scores: evidence from US school districts. Am Econ Rev Insights 5(2):173–90

Jones CJ, Philippon T, Venkateswaran V (2020) Optimal mitigation policies in a pandemic: social distancing and working from home. NBER Working Paper

Keppo J, Quercioli E, Kudlyak M, Smith L, Wilson A (2020) The behavioral SIR model, with applications to the swine Flu and COVID-19 pandemics

Kofoed M, Gebhart L, Gilmore D, Moschitto R (2024) Zooming to Class?: Experimental Evidence on College Students' Online Learning during COVID-19. Insights, American Economic Review (Forthcoming)

Konda A, Prakash A, Moss GA, Schmoldt M, Grant GD, Guha S (2020) Aerosol Filtration efficiency of common fabrics used in respiratory cloth masks. ACS Nano 14(5):6339–6347

Ma Q-X, Shan H, Zhang H-L, Li G-M, Yang R-M, Chen J-M (2020) Potential utilities of mask-wearing and instant hand hygiene for fighting SARS-CoV-2. J Med Virol 92(9):1567–1571

McAdams D (2021) The blossoming of economic epidemiology. Annu Rev Econ 13:539–570

McCormack G, Avery C, Spitzer AK-L, Chandra A (2020) Economic vulnerability of households with essential workers. JAMA 324(4):388–390

Oster E (2020) Schools aren't super-spreaders. *The Atlantic*. Last modified October 9, 2020. Retrieved August 23, 2024 from https://www.theatlantic.com/ideas/archive/2020/10/schools-arent-superspreaders/616669/

Osterholm MT, Olshaker M (2020) Deadliest enemy: our war against killer germs. Little, Brown and Company

Rachel L (2024) The worst-case scenario: optimal and equilibrium epidemic mitigation with moderately effective and fiscally costly lockdowns. Review of Economic Design. Forthcoming

Romer P (2020a) Simulating Covid-19: Part 1. March 23. Available at https://paulromer.net/covid-sim-part1/

Romer P (2020b) Simulating Covid-19: part 2. March 24. Available at https://paulromer.net/covid-sim-part2/

Romer P (2020c) Even a bad test can help guide the decision to isolate: covid simulations part 3. March 25. Available at https://paulromer.net/covid-sim-part3/

Roth AE (2002) The economist as engineer: game theory, experimentation, and computation as tools for design economics. Econometrica 70(4):1341–1378

Stock JH (2020) Coronavirus data gaps and the policy response to the novel coronavirus. NBER Working paper 26902

Stock JH, Aspelund KM, Droste M, Walker CD (2020) Estimates of the undetected rate among the SARS-CoV-2 infected using testing data from iceland. *medRxiv*

Taipale J, Romer P, Linnarsson S (2020) Population-scale testing can suppress the spread of COVID-19. medRxiv

Toxvaerd, Flavio. (2020). "Equilibrium Social Distancing." Cambridge Working Paper in Economics

Varma JK, Feldkamp C, Alexander M, Norman E, Agerton T, Davis R, Long T (2022) COVID-19 transmission due to delta variant in New York City public schools from october to december 2021. JAMA Netw Open 5(5):e2213276–e2213276

Varma JK, Thamkittikasem J, Whittemore K, Alexander M, Stephens DH, Arslanian K, Bray J, Long TG (2021) COVID-19 infections among students and staff in New York City public schools. Pediatrics 147(5):e2021050605

Yamey G (2021) Rich countries should tithe their vaccines. Nature 590:529