Social Distancing Equilibrium

Ruggiero Cavallo
Yahoo Research
New York, NY
cavallo@yahoo-inc.com

February 26, 2021

Abstract

In a pandemic people naturally modulate their social interactions, weighing rewards against risks, which are at least partly determined by the prevailing contact level of the broader community. If all are optimizing with respect to the decisions of others, what happens in equilibrium? In a highly simplified model, I highlight the collective action problem and find that: the lowest-contact equilibrium is always the best from a social welfare perspective, but still involves contact beyond the optimal rate; a small increase in viral transmissibility can do dramatically disproportionate damage to equilibrium welfare; and, on the other hand, facilitating a marginally greater social distancing option (think of curbside pickup, work from home, etc.) can sometimes eliminate the worst equilibria. The same is true for certain manipulative interventions, which raises questions about the communication strategies public health authorities may be drawn to engage in pursuit of social welfare.

1 Introduction

To social distance, or not to social distance? In a pandemic it becomes a pervasive question, assuming a society free enough to let it arise. Should you go to the barber, or trim your own hair; go out for dinner, or cook at home; catch up with a friend in person, or over the phone? These choices involve weighing benefit against cost, reward against risk; and the risk side of the equation is strongly impacted by the background decisions of the broader population.

Navigating this terrain involves game theoretic reasoning, with individual strategies optimized and tuned against the choices of others. Laxity of the broader society may provide more reason for you to be cautious, and inversely a highly cautious society can provide more scope for you to be lax without harm. Or does that logic flip in certain cases? What outcomes obtain? In other words, what is the *social distancing equilibrium*?

The unmeasurable complexities of the real world have confounded myriad attempts to predict the current pandemic's course, and this paper should not be confused for any such effort. Rather, I adopt a purposefully simplified model and focus on an under-appreciated dimension, the strategic one, revealing a number of salient patterns which I hope can beneficially enhance reasoning in this domain.

1.1 Main questions and answers

The central contribution is to model how equilibrium distancing levels might: 1) relate to social optima, and 2) change as a function of key variables (i.e., costs of infection, transmissibility rate, maximum amount of distancing that is possible, etc.).

Regarding social optima, there will be a free-riding problem in some form — when an individual decides hows much to socially distance, he is trading off a good (value from contacts) against a bad (possibility of infection); he receives the entire good himself,[1] but only experiences *part* of the bad, since his becoming infected may have serious negative ramifications on others in addition to the obvious personal cost. Indeed, I find that under my model, aside from a corner case:

> *Every equilibrium involves more social contact than occurs in the social optimum (Proposition 1).*

The second finding about welfare is perhaps slightly less intuitive:

[1] This is one of the over-simplifications of the model I propose, since social interaction is multi-way and in reality often brings some value to all sides. If you're riding the bus or getting groceries, you may bring no additional value to other passengers or shoppers; but if you're performing in a band, it's a different story.

The equilibrium generating highest social welfare is always the one that involves the least social contact (Proposition 2).

A revealing pattern, illustrated in Figure 1, is sometimes seen in individual best-response as a function of the social distancing decisions of others, with unmitigated contact optimal at *both* extremes (i.e., when the broader society is maximally distanced *and* when it is minimally distanced), and, in the middle, optimality of progressively greater distancing in response to diminishing forbearance in the broader society. Unfortunately, in such cases there is an unrestricted contact ("give up on distancing") equilibrium that is often very bad from a welfare perspective; one focus of the paper will be to understand conditions that can yield this outcome.

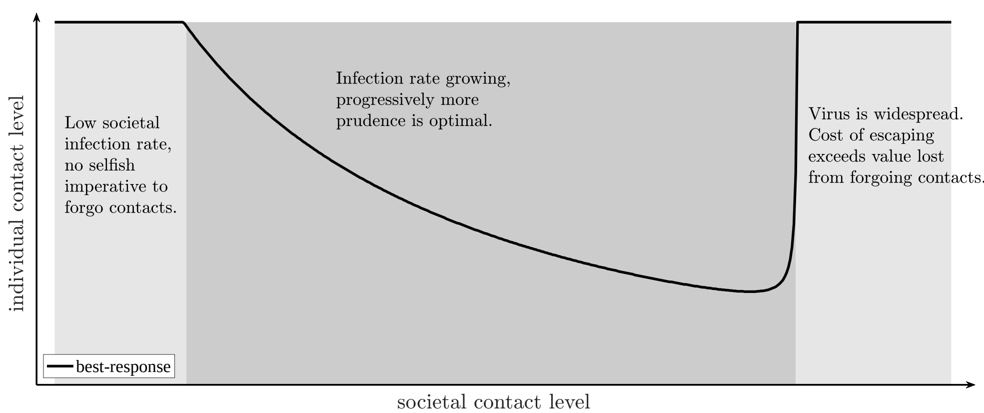

Figure 1: An illustrative pattern of best-response to others' distancing decisions. Unrestricted contact (zero social distancing) is one equilibrium here.

With respect to the second class of questions, regarding how outcomes change in response to input changes, I focus on what an increase in viral *transmissibility* portends, and also on the types of interventions we might expect from public health authorities. I find that, though in many cases increased transmissibility will just smoothly increase the equilibrium amount of social distancing, also:

An increase in transmissibility can abruptly introduce zero social distancing as an equilibrium; and thus a small increase in transmissibility can do vastly disproportionate harm to social welfare (Claim 1).

Regarding possible interventions of a social planner, we observe that in the current pandemic key authorities have, at least at times, exhibited a disposition to using their communication platforms as ways to *nudge and optimize* the public's behavior, rather than merely as channels through which to forthrightly inform.

For instance, when asked in a congressional hearing [10] about his failure to recommend mask-wearing amongst the general public throughout the early stage of the COVID-19 pandemic,[2] Dr. Anthony Fauci—the de facto spokesman for US public health policy at the time—explained that "there was a paucity of equipment that our heath care providers needed.... Now that we have enough, we recommend [widespread mask wearing]." In an interview with the *New York Times* [9], Fauci also described a strategic approach to communicating thresholds for herd immunity — tuning his estimates to how he expected the public to respond, rather than just straightforwardly conveying current expert opinion.[3]

Such examples motivate an exploration of potential informational "optimizations" (manipulations, really) that a welfare-minded authority might engage. I find:

> *A zero social distancing equilibrium can always be eliminated by sufficiently elevating the perceived relative "cost" of infection as compared to the value of social contact (Theorem 3). In some cases a tiny movement along these lines will change the equilibrium outcome from terrible to nearly optimal.*

This should raise an alarm about some of the temptations authorities might face, focused on optimizing a set of key metrics, without good visibility into the specific and often dire tradeoffs faced by individuals ("the best-laid schemes o' mice an' men").

Finally, there may be forthright interventions that also have a remarkably salutary effect on social welfare. In particular:

---

[2] "There's no reason to be walking around with a mask," Fauci told *60 Minutes* in an interview that aired March 8, 2020 (`https://www.cbsnews.com/news/preventing-coronavirus-facemask-60-minutes-2020-03-08`).

[3] Quoting from the *New York Times*: *"When polls said only about half of all Americans would take a vaccine, I was saying herd immunity would take 70 to 75 percent," Dr. Fauci said. "Then, when newer surveys said 60 percent or more would take it, I thought, 'I can nudge this up a bit,' so I went to 80, 85."*

> *A disastrous zero social distancing equilibrium can sometimes be eliminated by marginally increasing the level of social distancing that is possible (Theorem 4).*

Intuitively, if the greatest possible social distancing level is *not distant enough* to make it worth the cost of forgoing valuable contact, then the only equilibrium may be zero distancing rather than the maximum or some intermediate level. Making it feasible to engage in lower contact levels—think of grocery delivery, curbside pickup, work from home, etc.—can create a far better equilibrium.

1.2 Limitations

I adopt a highly simplified model with many assumptions that deviate starkly from reality. I assume a *uniform population* and analyze symmetric outcomes, whereas in reality there are vast differences in: cost of infection across individuals (e.g., old versus young), ability to social distance (e.g., essential on-site worker versus home worker), risk imposed by one potential contact versus another (e.g., close co-worker versus clerk), etc.

I assume that agents perfectly optimize their utilities, while in reality many are uninformed and all are vulnerable to error.

I assume agents are *selfish*, while in reality, one hopes, most are highly motivated to protect the well-being of others.

I ignore the fact that the space of social distancing options available across agents may be constrained interdependently.[4]

Also, I model optimization with respect to any given "snapshot" of the world, whereas in reality things are constantly in flux (and agents know this) — new vaccines on the horizon, new viral strains evolving, etc. I cannot report on how robust my results would be, qualitatively, to de-simplification along these dimensions, if such is possible.

Finally, while motivated by the COVID-19 pandemic, my analysis does not make any attempt to find a set of model parameters that approximates the current situation, or any specific historical situation for that matter. Instead I seek general insights, and derive findings about what interesting things are *possible* across the space of potential scenarios.

[4] E.g., I assume "no social distancing" is an option even when all others are maximally distancing, while in reality there is no way of maintaining the same level of social contact when, say, the subways are only 30% as full as before. (This is related to footnote 1.)

1.3 Related work

A fundamental premise underlying this work is that people will naturally alter their behavior, distancing from others, when faced with a dangerous contagion. Practically anyone alive during the current global pandemic has seen this first-hand, but studies such as [1] catalog evidence that even in a far more subtle case such as the 2009 H1N1 flu, Americans spontaneously social distanced in a way that impacted the epidemic.[5]

Other work makes the case for doing so. Ferguson et al. [7] estimated that in the "unmitigated" case, COVID-19 would lead to over 2 million US deaths by the end of Summer 2020, with perhaps *half* avoidable through social distancing and targeted isolation measures.[6] The quantitative predictions of this study and its derivatives[7] are not easy to evaluate, since a mix of mitigation and more stringent suppression interventions were ultimately adopted, but they at least provide an indication of how critical distancing is seen to be by those most prominent in the field.

With respect to game theoretic modeling, prior work contextualizes a social distancing "game" within the canonical SIR (susceptible, infected, recovered) approach to epidemiological modeling. Reluga [14] models how behaviors might change over time as, say, a vaccination horizon draws nearer. Chen et al. [4] models the equilibrium distancing rate as a function of how the amount of social contact relates to the number of people "out in public." Like some others, Farboodi et al. [5] compares outcomes under a laissez-faire policy and an optimal policy that imposes distancing constraints, but is distinguished by its specific application to the COVID-19 crisis. Fenichel [6] emphasizes the importance of discriminating amongst different health classes when formulating distancing policies, and considers agents with value that is concave in their chosen contact rate (as do I). Finally, Toxvaerd [16] considers settings with a continuum of agents, which simplifies the equilibrium analysis (some related assumptions I make in this work play a similar role, though we get qualitatively different results). Chang [3] provides a broad review of work in this area.

[5]Interestingly, the authors argue that the under-appreciation of voluntary distancing has in the past lead to excessively strict formal measures.

[6]The computer simulation code underlying this report's modeling and predictions has since received intense criticism from some quarters; see [15].

[7]Using Ferguson et al.'s simulation model, Greenstone and Nigam [8] estimated that "3-4 months of moderate distancing beginning in late March 2020 would save 1.7 million lives in the US by October 1."

I do something quite different in this paper, bypassing the SIR model (which adds a lot of complexity while, at least in this context, delivering questionable predictive value), instead collapsing the time dimension and simply including the most game-theoretically relevant datum—namely, how individual infection risk varies with the distancing level of others—as a *direct input* the modeler can specify. This yields findings that are highly intelligible, since they directly describe relations of fundamental and quasi-observable attributes of the problem.

Finally, at the end of the paper I discuss manipulations a welfare-minded planner might employ to try to mitigate the free-rider dynamic at the core of collective action problems; this relates very directly to [2].

2 Model

A large population of agents simultaneously decide on individual *contact levels*, jointly determining the utility experienced by each, which is a linear combination of *value* from the social interaction and *cost* of infection risk.

I will often present the scenario from the perspective of an individual agent i, denoting his chosen contact level as $c_i \in [\underline{c}, 1]$ and that of the others as $c_o \in [\underline{c}, 1]$ (for a $\underline{c}$ to be described below). Reducing the contact level of "others" to a scalar is a simplification that glosses over any strategic distinctions between, say, a scenario where half the population locks down and half is unconstrained, versus one where everyone distances at some intermediate rate; in this model, agents optimize with respect to a thus-generalized measure of societal contact. I will formalize this further in a moment.

Each agent has m potential contacts, and a choice of c_i means he will interact with $c_i m$ of them. $\underline{c} > 0$ captures the fact that total isolation is impossible (e.g., $\underline{c} = \frac{1}{m}$ would indicate that all must make at least 1 contact).

Due to the dynamics of contagion, the probability that a random contact is infected is increasing in the contact rate of the population. Thus i's probability of becoming infected from a *single* contact is an increasing function f of c_o.[8] I assume only that this function f is continuous, increasing, and

[8]One intuitive form for f would be $f(c_o) = \beta[r + (1-r)c_o^\sigma]$ with parameters $\beta \in (0,1)$, $r \in (0,1)$, and $\sigma \geq 1$ representing likelihood of transmission given contact with an infected person, infection *base rate* in the population, and a slope-control, respectively. We will return to a specific form like this for demonstration purposes later in the paper, but our results can be stated generally.

has range $\in [0,1]$ for all $c_o \in [\underline{c}, 1]$. I will use $\delta(c_o)$ to denote an uninfected agent's probability of *remaining uninfected* given a single contact interaction; i.e., $\delta(c_o)$ is shorthand for $1 - f(c_o)$. Then i's total probability of becoming infected given c_i (i.e., given $c_i m$ contacts), equals $1 - \delta(c_o)^{c_i m}$.[9] Note that δ is continuous, decreasing, and has range $\in [0,1]$ for all $c_o \in [\underline{c}, 1]$.

Now returning to c_o, it can be interpreted as whatever aggregate measure of societal distancing determines the single-contact non-infection probability, δ, as perceived by agents — it could be a statistic like the mean or mode over individual contact levels, or something else; there is no need to pin it down. Formally, I assume only that all c_o values in $[\underline{c}, 1]$ are realizable (i.e., obtain for some set of individual choices), and in the case that every i has chosen $c_i = x$, for arbitrary $x \in [\underline{c}, 1]$, that $c_o = x$ for each i.

On the other side of the ledger is the *value* i gets from interaction, which exhibits diminishing marginal returns as contacts are added.[10] I represent this as $\alpha(1 - \gamma^{c_i m})$, where $\gamma \in (0,1)$ is a parameter controlling the rate of exponential decay (e.g., if $\gamma = 0.75$, each additional contact is three quarters as valuable as the previous one), and $\alpha > 0$ is a way of scaling value against the cost of infection.

Putting this together and subtracting cost (infection probability) from value yields the following **utility function**:

$$u_i(c_i, c_o) = \Big[\alpha(1 - \gamma^{c_i m})\Big] - \Big[1 - \delta(c_o)^{c_i m}\Big] \tag{1}$$

This implicitly normalizes things to a "cost of infection" that equals 1, with the relative "contact value" scalable via α without changing the rate of decay.

In sum, c_i and c_o are variables that agents jointly (but individually) set, while scale parameter α, value discount rate γ, minimal contact rate

[9]This expression for infection probability makes sense regardless of the time-scale one envisions for "a contact." Your (perhaps repeated) interaction with a single contact, say, the grocer, brings a probability of infection that will be distinct depending on whether the time-frame being considered is 1 day or 1 month, but either choice can be the basis for specifying input $\delta(c_o)$, and in either case the additional infection probability—*over the same time-frame*—for adding a second contact is: $[1 - \delta(c_o)^2] - [1 - \delta(c_o)]$.

[10]Contacts have varying levels of importance and distancing will naturally eliminate the least important first. Contact with the grocer is critical for getting food; then perhaps extended family contacts, while not strictly necessary and thus marginally less "valuable," are certainly still quite important. Then, perhaps the dentist, the butcher, and so on and so forth until we come to far less critical contact with strangers at the bottom.

$\underline{c}$, contact-pool size m, and the single-contact non-infection probability function $\delta : [\underline{c}, 1] \to [0, 1]$, are inputs included as variables to capture a breadth of scenarios. An example is depicted in Figure 2.

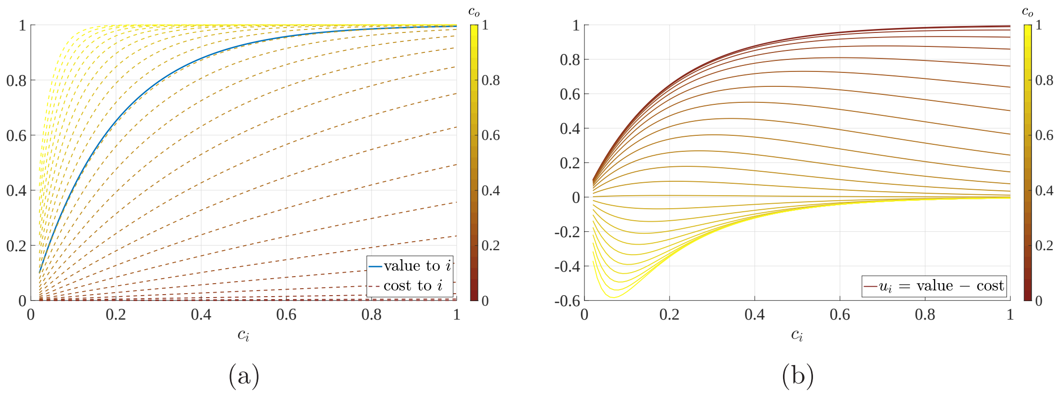

Figure 2: For parameters $m = 50$, $\alpha = 1$, $\gamma = 0.9$, and $\delta(c_o) = 1 - 0.5c_o^3$, value and costs (a) and utility (b) for agent i are illustrated as a function of his chosen contact level (c_i, on the x-axes) and that chosen by others (c_o, distinguished by color). Plot lines start at the minimal feasible contact level, $\underline{c} = 0.02$.

3 Best-response

We begin analysis in this model by mapping out the optimal contact level choices for a given agent i, conditional on whatever contact level c_o is operative in the broader society.

The function below, which we will see relates to the first-order conditions for utility maximization, plays an important role in what follows. Define:

$$\chi(c_o) = \frac{\ln\left(\frac{\alpha \ln(\gamma)}{\ln(\delta(c_o))}\right)}{m \ln\left(\frac{\delta(c_o)}{\gamma}\right)} \tag{2}$$

There are 3 cases to address, distinguished by how the single-contact non-infection probability (δ) relates to the rate at which marginal value decays as contacts increase (γ). The best-response is either $\chi(c_o)$, or it lies at one of the extremes ($\underline{c}$ or 1).

Lemma 1. *Given c_o, i's best-responses are as follows:*

- *If $\delta(c_o) > \gamma$, i has a unique best-response:* $\min\{1, \max\{\underline{c}, \chi(c_o)\}\}$.
- *If $\delta(c_o) < \gamma$, i has a unique best-response $\underline{c}$ if $\alpha(\gamma^{\underline{c}m} - \gamma^m) < \delta(c_o)^{\underline{c}m} - \delta(c_o)^m$, a unique best-response 1 if $\alpha(\gamma^{\underline{c}m} - \gamma^m) > \delta(c_o)^{\underline{c}m} - \delta(c_o)^m$, and two best-responses $\underline{c}$ and 1 otherwise.*
- *If $\delta(c_o) = \gamma$: i has a unique best-response $\underline{c}$ if $\alpha < 1$, i has a unique best-response 1 if $\alpha > 1$, and every $c_i \in [\underline{c}, 1]$ is a best-response if $\alpha = 1$.*

Proof. A best-response for i to c_o is a global optimum of $u_i(c_i, c_o)$, fixing c_o. In the corner case where $c_o = \delta^{-1}(\gamma)$, i.e., $\delta(c_o) = \gamma$, $u_i(c_i, c_o) = (\alpha - 1)(1 - \gamma^{c_i})$, and $\frac{\partial}{\partial c_i} u_i(c_i, c_o) = (1 - \alpha)\ln(\gamma)\gamma^{c_i}$. If $\alpha < 1$ this is negative everywhere, and so $c_i = \underline{c}$ is optimal; if $\alpha > 1$ this is positive everywhere, and so $c_i = 1$ is optimal; and if $\alpha = 1$ this is 0 everywhere, and so every c_i is optimal.

For the rest of the proof, we assume $c_o \neq \delta^{-1}(\gamma)$. We have:

$$\frac{\partial}{\partial c_i} u_i(c_i, c_o) = m \ln(\delta(c_o))\delta(c_o)^{c_i m} - \alpha m \ln(\gamma)\gamma^{c_i m}$$

We are interested in the sign of this quantity, and whether and where the sign changes. $\frac{\partial}{\partial c_i} u_i(c_i, c_o) > 0$ if and only if $\ln(\delta(c_o))\delta(c_o)^{c_i m} > \alpha \ln(\gamma)\gamma^{c_i m}$, i.e., noting that $\ln(\delta(c_o)) < 0$, $\frac{\partial}{\partial c_i} u_i(c_i, c_o) > 0$ if and only if:

$$\left(\frac{\delta(c_o)}{\gamma}\right)^{c_i m} < \frac{\alpha \ln(\gamma)}{\ln(\delta(c_o))} \tag{3}$$

$\chi(c_o)$, defined above in Eq. (2), is precisely the c_i value that equalizes both sides of this inequality.

Now we need to distinguish two sets of cases, depending on whether or not $\delta(c_o) > \gamma$. First assume it is so, i.e., that $c_o < \delta^{-1}(\gamma)$. Eq. (3) will be satisfied, and thus $\frac{\partial}{\partial c_i} u_i(c_i, c_o)$ will be positive, if and only if $c_i < \chi(c_o)$. That means u_i will be increasing for all $c_i < \chi(c_o)$ and decreasing for all $c_i > \chi(c_o)$, and thus the best-response is $\underline{c}$ if $\chi(c_o) \leq \underline{c}$, $\chi(c_o)$ if $\chi(c_o) \in [\underline{c}, 1]$, and 1 if $\chi(c_o) > 1$. In other words, the best-response is $\min\{1, \max\{\underline{c}, \chi(c_o)\}\}$.

Now assume $\delta(c_o) < \gamma$, i.e., $c_o > \delta^{-1}(\gamma)$. This time (3) will be satisfied, and $\frac{\partial}{\partial c_i} u_i(c_i, c_o)$ will be positive, if and only if $c_i > \chi(c_o)$. That means u_i will be decreasing for all $c_i < \chi(c_o)$ and increasing for all $c_i > \chi(c_o)$. The best-response must therefore lie at an extreme — it is $\underline{c}$ if $\alpha(\gamma^{\underline{c}m} - \gamma^m) < \delta(c_o)^{\underline{c}m} - \delta(c_o)^m$, 1 if $\alpha(\gamma^{\underline{c}m} - \gamma^m) > \delta(c_o)^{\underline{c}m} - \delta(c_o)^m$, and both extremes are best-responses otherwise (we know that the first case must obtain if $\chi(c_o) \geq 1$). $\square$

Corollary 1. *For every $c_o \in [\underline{c}, 1]$, for all possible parameter settings with either $\delta(c_o) \neq \gamma$ or $\alpha \neq 1$, i's best-response c_i is either $\underline{c}$, 1, or $\chi(c_o) = \ln\left(\frac{\alpha \ln(\gamma)}{\ln(\delta(c_o))}\right) / \left[m \ln\left(\frac{\delta(c_o)}{\gamma}\right)\right]$. The best-response is unique except when $\delta(c_o) < \gamma$ and $\alpha(\gamma^m - \gamma^{\underline{c}m}) = \delta(c_o)^m - \delta(c_o)^{\underline{c}m}$, in which case $\underline{c}$ and 1 are both best-responses, or when $\delta(c_o) = \gamma$ and $\alpha = 1$, in which case all c_i are best-responses.*

Thus best-response is essentially a function from c_o to c_i, and we can map out its form. Figure 3 illustrates one example, an intuitive case where, when others' contact levels are minimal and infection risk is therefore relatively low, no distancing ($c_i = 1$) is an optimal response; as others' contact levels increase, some distancing becomes optimal, at an increasing level as c_o rises, until eventually when others are very unconstrained minimal contact ($c_i = \underline{c}$) is a best-response.

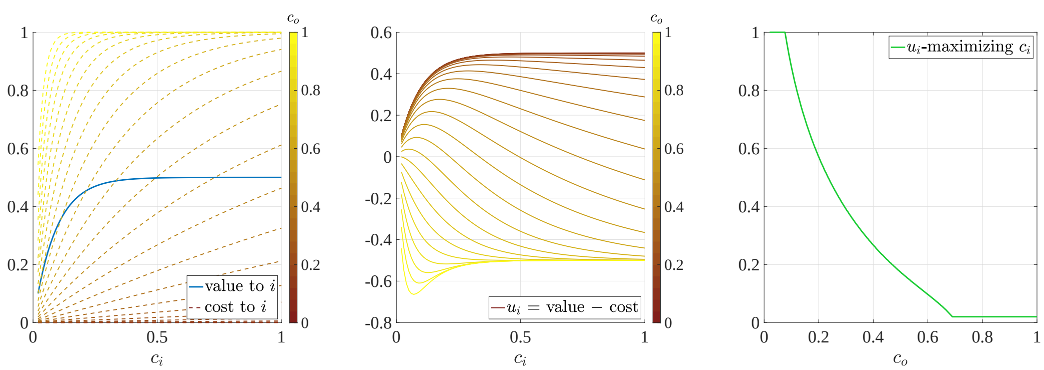

(a) (b) (c)

Figure 3: Example value and costs (a), and utility (b), with the resulting best-response (c) as a function of the contact level c_o chosen by others. For parameters $m = 50$, $\alpha = 0.5$, $\gamma = 0.8$, and $\delta(c_o) = 1 - 0.6c_o^5$.

3.1 Equilibrium

Consider a large overall population, with a subpopulation of m potential contacts per agent, and imagine that each agent's individual influence on any other's utility function is minute in expectation, e.g., as would be the case if the m potential contacts are drawn quasi-randomly (think of a random collection of fellow diners, passengers on a bus, etc., on a given day).[11]

If no *single* agent meaningfully impacted any other's expected utility, and if best-responses were unique (and we've seen that they are, modulo the very small number of corner cases spelled out in Lemma 1), then any possible equilibrium would necessarily be *symmetric*, i.e., with $c_i = c_o$ for every agent i. The "large population" assumption means all agents face the same effective c_o; that combined with uniqueness of best-response means they will all have the same optimal reaction.

With that informal motivation, we will restrict our attention to symmetric outcomes going forward: I will use the term *equilibrium* to mean *symmetric pure strategy Nash equilibrium*, and accordingly phrases like "there exists no equilibrium" should not be taken to exclude the possible existence of other types of equilibria.

I will henceforth use $b(c_o)$ to denote the best-response function, characterized in Lemma 1. An equilibrium occurs at any c that has the property of being an element of $b(c)$. Visually, this corresponds to an intersection between the best-response curve and a 45-degree line (see Figure 4).

Theorem 1. *Equilibrium characterization.*

- $\underline{c}$ *is an equilibrium if and only if:*

$$\begin{aligned} &\big(\delta(\underline{c}) > \gamma \ \wedge \ \chi(\underline{c}) \leq \underline{c}\big) \\ \vee \ &\big(\delta(\underline{c}) = \gamma \ \wedge \ \alpha \leq 1\big) \\ \vee \ &\big(\delta(\underline{c}) < \gamma \ \wedge \ \alpha(\gamma^{\underline{c}m} - \gamma^m) \leq \delta(\underline{c})^{\underline{c}m} - \delta(\underline{c})^m\big) \end{aligned}$$

- $c \in (\underline{c}, 1)$ *is an equilibrium if and only if:*

$$\big(\delta(c) > \gamma \ \wedge \ \chi(c) = c\big) \ \vee \ \big(\delta(c) = \gamma \ \wedge \ \alpha = 1\big)$$

[11] Admittedly, this completely fails to capture important real-world asymmetries: within-household contacts are vastly more relevant than one-time encounters with passengers on a train, etc.

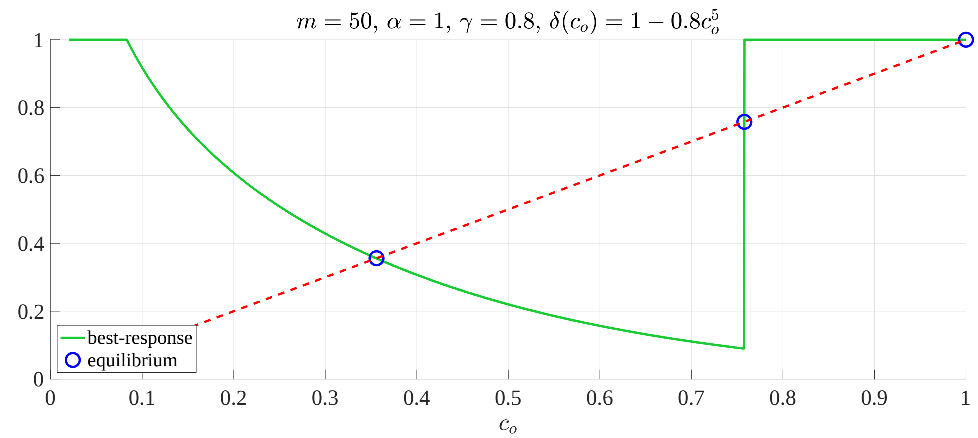

Figure 4: Equilibria correspond to the points where the best-response curve intersects a 45-degree line from the origin. In this case there are three: one at a point $c \approx 0.35$ where $\delta(c) > \gamma$, one where $\delta(c) = \gamma$ (since $\alpha = 1$ here), and one at $c = 1$ where $\delta(1) < \gamma$.

- 1 *is an equilibrium if and only if:*

$$\begin{aligned} &\big(\delta(1) > \gamma \ \wedge \ \chi(1) \geq 1\big) \\ \vee\ &\big(\delta(1) = \gamma \ \wedge \ \alpha \geq 1\big) \\ \vee\ &\big(\delta(1) < \gamma \ \wedge \ \alpha(\gamma^{\underline{c}m} - \gamma^m) \geq \delta(1)^{\underline{c}m} - \delta(1)^m\big) \end{aligned}$$

Proof. For every $c \in [\underline{c}, 1]$, c is an equilibrium if and only if $c \in b(c)$. Applying Lemma 1 to $c = \underline{c}$, $c \in (\underline{c}, 1)$, and $c = 1$ in turn, tracking the cases where $c \in b(c)$, directly yields the theorem. □

The following very mild condition is sufficient to guarantee the existence of at least one equilibrium.

Condition 1.

$$\alpha \notin \left[1, \frac{\delta(1)^{\underline{c}m} - \delta(1)^m}{\gamma^{\underline{c}m} - \gamma^m}\right]$$

This condition will frequently be satisfied trivially (i.e., necessarily), because $\delta(1)^{\underline{c}m} - \delta(1)^m$ may be less than $\gamma^{\underline{c}m} - \gamma^m$, yielding an empty interval.[12]

[12]For instance, say the number of possible contacts (m) is 50, and $\underline{c} = 0.02$ (i.e., the

Theorem 2. *Given Condition 1, there is always an equilibrium, and for all but at most two possible values of α, there is an equilibrium in which all agents are playing a unique best-response.*

Proof. Assume Condition 1. We need to show that there exists at least one $c_o \in [\underline{c}, 1]$ such that $c_o \in b(c_o)$, and *only* one such c_o for all but two possible values that α could take.

If $\delta(1) > \gamma$, then $b(c_o) = \{\min\{1, \max\{\underline{c}, \chi(c_o)\}\}\}$, $\forall c_o \in [\underline{c}, 1]$ (see Lemma 1), and since this is a continuous function with range $\in [0, 1]$, $b(c_o) = \{c_o\}$ at some $c_o \in [\underline{c}, 1]$.

If $\delta(1) \leq \gamma$ and $\alpha > \frac{\delta(1)^{\underline{c}m} - \delta(1)^m}{\gamma^{\underline{c}m} - \gamma^m}$, then $u_i(1, 1) > u_i(\underline{c}, 1)$ and thus $b(1) = \{1\}$.

If $\delta(1) < \delta(\underline{c}) \leq \gamma$ and $\alpha < 1$, if there were no equilibrium then necessarily $u_i(\underline{c}, \underline{c}) < u_i(1, \underline{c})$ and $u_i(\underline{c}, 1) > u_i(1, 1)$, i.e.,

$$\delta(\underline{c})^{\underline{c}m} - \delta(\underline{c})^m < \alpha(\gamma^{\underline{c}m} - \gamma^m) < \delta(1)^{\underline{c}m} - \delta(1)^m \tag{4}$$

Consider the function $h(k) = k^{\underline{c}m} - k^m$. As k goes from $\underline{c}$ to 1, $h'(k) = mk^{\underline{c}m-1}(1 - k^{m(1-\underline{c})})$ undergoes a single change in sign from positive to negative; i.e., $h(k)$ increases at first and then decreases. Thus Eq. (4) is incompatible with $\delta(1) < \delta(\underline{c}) \leq \gamma$, given that $\alpha < 1$,[13] and either $\underline{c} \in b(\underline{c})$ (uniquely so if $\alpha \neq [\delta(\underline{c})^{\underline{c}m} - \delta(\underline{c})^m]/[\gamma^{\underline{c}m} - \gamma^m]$) or $1 \in b(1)$ (uniquely so if $\alpha \neq [\delta(1)^{\underline{c}m} - \delta(1)^m]/[\gamma^{\underline{c}m} - \gamma^m]$).

Finally if $\delta(1) \leq \gamma < \delta(\underline{c})$ and $\alpha < 1$, then $\forall c_o \in [\underline{c}, \delta^{-1}(\gamma)]$, $b(c_o) = \min\{1, \max\{\underline{c}, \chi(c_o)\}\}$. But given that $\alpha < 1$, as c_o approaches $\delta^{-1}(\gamma)$ from the left, $\min\{1, \max\{\underline{c}, \chi(c_o)\}\}$ necessarily (smoothly) goes to and reaches $\underline{c}$. This holds because the numerator of Eq. (2) converges to a negative constant, $\ln(\alpha)$, while the denominator is non-negative and converges to 0. Thus $b(c_o) = \min\{1, \max\{\underline{c}, \chi(c_o)\}\}$ must equal c_0 at at least one point $c_o \in [\underline{c}, \delta^{-1}(\gamma))$. □

As the condition perhaps suggests, there are indeed cases outside its specifications that have no equilibrium. See Figure 5 for an example. In it, $\chi(c) > 1$ for all $c \in [\underline{c}, \delta^{-1}(\gamma))$, making maximal contact the best-response

minimal number of contacts is 1). If the probability of non-infection given a single contact, assuming that others adopt no social distancing ($\delta(1)$), is less than the value discount rate (γ), then for any $\gamma \leq 0.92$ the condition is guaranteed to be satisfied.

[13] $h(\delta(1)) > h(\delta(\underline{c}))$ entails that $\delta(\underline{c})$ is after the peak, but $h(\delta(\underline{c})) < \alpha h(\gamma)$ entails that it is before.

over that interval; then at a point prior to $c_o = 1$ the best-response switches from maximal contact to minimal contact, excluding the possibility of an equilibrium on the rest of the interval. Still, as Theorem 2 indicates, this is an odd exception to the rule.

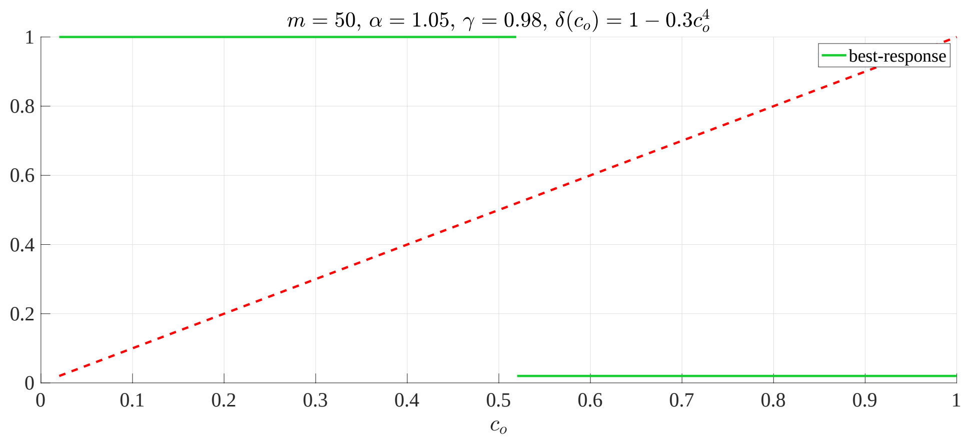

Figure 5: An example with no equilibrium. As others distance less, maximal-distancing replaces no-distancing as the best-response.

4 Social welfare

Now that we know *what* the equilibria are for a given set of inputs, we can evaluate *how good* they are. We will first of all see that the free-rider problem is indeed a problem here, for all inputs that yield anything other than only a minimal- or maximal-contact equilibrium (in the latter case there's nothing left to free-ride on).

Proposition 1. *In any equilibrium, the contact level is greater than or equal to a socially optimal level; if the socially optimal contact level is non-extremal, then any equilibrium contact level will be strictly greater than optimal.*

Proof. Let $\dot{c}$ be a socially optimal contact level, i.e., omitting some irrelevant constants from the utility function Eq. (1),[14]

$$\dot{c} \in \underset{c \in [\underline{c}, 1]}{\arg\max} \Big(\delta(c)^{cm} - \alpha\gamma^{cm}\Big)$$

If $c \in [\underline{c}, 1]$ is an equilibrium, then all agents are best-responding, and therefore:

$$\delta(c)^{cm} - \alpha\gamma^{cm} \geq \delta(c)^{\dot{c}m} - \alpha\gamma^{\dot{c}m}$$

But $c < \dot{c}$ would entail (recalling that δ is decreasing):

$$\delta(c)^{\dot{c}m} - \alpha\gamma^{\dot{c}m} > \delta(\dot{c})^{\dot{c}m} - \alpha\gamma^{\dot{c}m}$$

This contradicts $\delta(c)^{cm} - \alpha\gamma^{cm} \leq \delta(\dot{c})^{\dot{c}m} - \alpha\gamma^{\dot{c}m}$, which must hold by optimality of $\dot{c}$. Therefore it must be the case that $c \geq \dot{c}$.

Given this, to prove the second half of the proposition it is sufficient to show that if $\dot{c} \in (\underline{c}, 1)$, then $\dot{c}$ is not an equilibrium. Letting $\mathrm{W}(c)$ denote the per-agent welfare when the (symmetric) contact level is c, we have $\mathrm{W}(c) = \delta(c)^{cm} - \alpha\gamma^{cm} + (\alpha - 1)$, and:

$$\mathrm{W}'(c) = \delta(c)^{cm}\Big(m\ln(\delta(c)) + \frac{mc\delta'(c)}{\delta(c)}\Big) - \alpha m \ln(\gamma)\gamma^{cm}$$

Therefore $\mathrm{W}'(c)\big|_{c=\dot{c}} = 0$, which is a necessary condition for optimality of non-extremal c, if and only if:

$$\delta(\dot{c})^{\dot{c}m} = \frac{\alpha\ln(\gamma)\gamma^{\dot{c}}m}{\ln(\delta(\dot{c})) + \frac{\dot{c}\delta'(\dot{c})}{\delta(\dot{c})}}$$

If the same $\dot{c}$ is an equilibrium, then the first-order conditions for the agent utility function must likewise be satisfied at $\dot{c}$. We have:

$$\frac{\partial}{\partial c_i} u_i(c_i, c) = m\Big(\ln(\delta(c))\delta(c)^{c_i m} - \alpha\ln(\gamma)\gamma^{c_i m}\Big)$$

And $\frac{\partial}{\partial c_i} u_i(c_i, c)\big|_{c_i = c = \dot{c}} = 0$ if and only if:

$$\ln(\delta(\dot{c}))\bigg(\frac{\alpha\ln(\gamma)\gamma^{\dot{c}m}}{\ln(\delta(\dot{c})) + \frac{\dot{c}\delta'(\dot{c})}{\delta(\dot{c})}}\bigg) - \alpha\ln(\gamma)\gamma^{c_i m} = 0$$

[14]Specifically, when all agents choose contact level c, all obtain the same welfare, which is $\delta(c)^{cm} - \alpha\gamma^{cm} + (\alpha - 1)$ per agent. It is a monotone transformation to omit the additive constant $\alpha - 1$.

I.e., if and only if:

$$\alpha \ln(\gamma)\gamma^{\dot{c}m}\left(\frac{\ln(\delta(\dot{c}))}{\ln(\delta(\dot{c})) + \frac{\dot{c}\delta'(\dot{c})}{\delta(\dot{c})}} - 1\right) = 0$$

This cannot be satisfied because δ is strictly decreasing everywhere on the interval, and the result follows. □

Not only are all equilibria "to the right of" (higher contact than) the optimum, but even if social-welfare dips and rises as the overall contact level increases, the best equilibrium is always the one involving least contact.

Proposition 2. *If there are multiple equilibria, the one with lowest contact level yields strictly greatest social welfare.*

Proof. Consider two arbitrary equilibria $c, \hat{c} \in [\underline{c}, 1]$ with $c < \hat{c}$. The fact that c is an equilibrium means there is no profitable deviation, and in particular $\hat{c}$ is not a profitable deviation. Moreover, $\delta(c) > \delta(\hat{c})$, since δ is monotonically decreasing. We thus have:

$$\begin{aligned} \delta(c)^{cm} - \alpha\gamma^{cm} &\geq \delta(c)^{\hat{c}m} - \alpha\gamma^{\hat{c}m} \\ &> \delta(\hat{c})^{\hat{c}m} - \alpha\gamma^{\hat{c}m} \end{aligned}$$

The first and last quantities are expressions of (a monotone simplification of) social welfare under c and $\hat{c}$, respectively. The claim follows. □

Figure 6 adds a plot of social welfare (assuming symmetric adoption of contact level c on the x-axis) to the plots from Figure 4. We can see how the example conforms to the above two propositions, with all three equilibria to the right of the social optimum, which occurs at $c \approx 0.287$ generating per-agent welfare ≈ 0.937, with the lowest equilibrium ($c \approx 0.356$, for per-agent welfare ≈ 0.903) generating vastly more welfare than the other two (which both yield welfare ≈ 0).

5 Transmissibility changes

Viruses evolve over time, and in some cases mutations can occur that yield variants with increased transmissibility or virulence. One such case is the emergence of a SARS-CoV-2 variant in the United Kingdom towards the

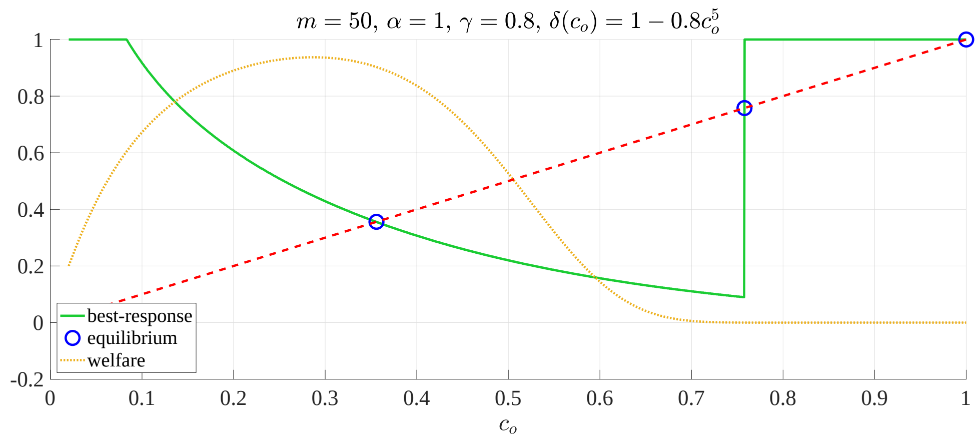

Figure 6: Welfare is plotted as a function of a universally adopted contact level. All equilibria are to the right of the welfare peak, and that yielding highest welfare is the one with lowest contact level.

end of 2020 that early studies estimated to be between 30% to 80% more transmissible than wild-type SARS-CoV-2 [12, 11]. If a more transmissible variant becomes predominant, the strategic considerations are changed.

To examine the possible impact, we can add one detail to δ, the function representing probability of remaining uninfected after a single contact. We imagine now that $\delta(c_o) = 1 - \beta g(c_o)$, for some $\beta \in (0, 1]$ that captures *the probability of transmission given that the contact is infected*, with g an arbitrary increasing function with range $(0, 1)$ representing the probability that a single contact will be an infected person, given c_o.

Considering the maximal-contact equilibrium condition in Theorem 1, noting that an increase in β can only yield a *decrease* in $\delta(c_o)$, we have the following informal claim that will apply very generally, if not in every specific case:

Claim 1. *An increase in transmissibility can introduce a maximal-contact equilibrium, but will not eliminate one.*

Informally, the first case of the maximal-contact equilibrium condition in Theorem 1 cannot be satisfied for any values of interest ($\delta(1)$, γ, and α would all have to be placed at wild extremes), and the second case is an extreme corner case that can be ignored for all practical purposes. That leaves the third condition, $\delta(1) < \gamma \ \wedge \ \alpha(\gamma^{\underline{c}m} - \gamma^m) \geq \delta(1)^{\underline{c}m} - \delta(1)^m$. If the first conjunct is satisfied for some value of β, then it clearly remains satisfied for all greater values. The same will hold for the second conjunct if we assume $\underline{c}m$ is relatively small, and m is large enough, because then the condition's satisfaction will hinge on whether $\alpha\gamma^{\underline{c}m} \geq \delta(1)^{\underline{c}m}$. As β increases, $\delta(1)$ decreases, only making the condition more readily satisfied.

Figure 7 illustrates this effect for parameters $m = 50$, $\underline{c} = 0.02$, $\gamma = 0.85$, $\alpha = 1$, and $\delta(c_o) = 1 - \beta c_o^{1.1}$. When β is quite small — i.e., when it is relatively unlikely that a single interaction with an infected person will yield transmission — the maximal-contact outcome ($c = 1$) is not an equilibrium. Once β reaches about 0.15, it is an equilibrium, and remains one for all greater values of β, with increasing "robustness" in the sense that maximal contact remains a best-response for an ever-widening range of choices by others.

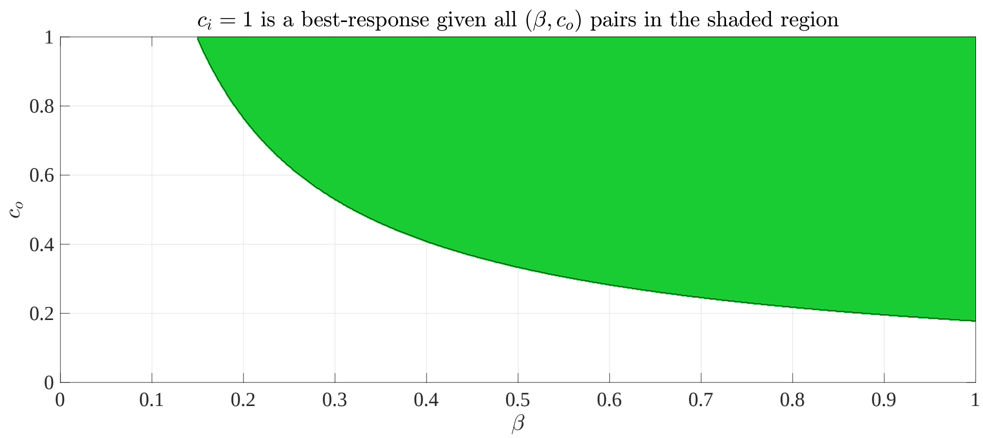

Figure 7: Illustration of how $c_i = 1$ becomes a best-response to an ever-broadening range of c_o as β grows, given $m = 50$, $\underline{c} = 0.02$, $\gamma = 0.85$, $\alpha = 1$, and $\delta(c_o) = 1 - \beta c_o^{1.1}$. $c = 1$ is thus an equilibrium for all and only $\beta > 0.15$.

Figure 8 illustrates equilibrium outcomes for two β values right around the threshold that makes maximal contact an equilibrium in the example.

6 Interventions, manipulative and otherwise

In a scenario with equilibria that are far from welfare-optimal, a social planner will naturally seek interventions to improve the situation, or ensure that

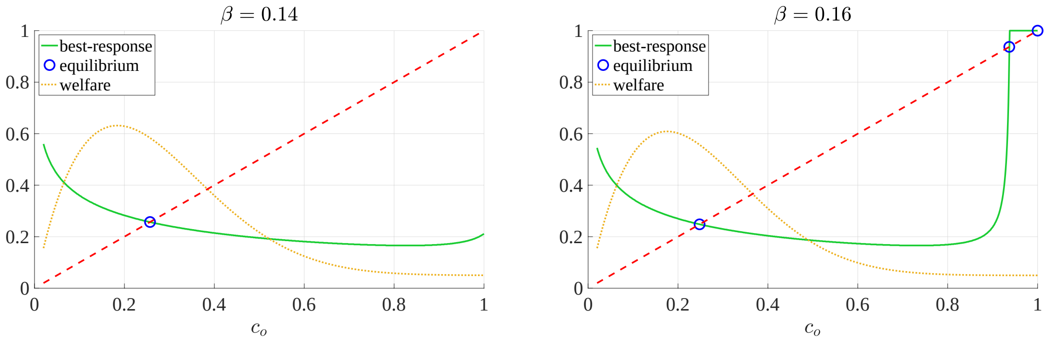

(a) (b)

Figure 8: A small increase in transmissibility (β) here transforms a scenario with a single high-welfare equilibrium (a) to one with three equilibria, two of which are essentially worst-case welfare outcomes (b).

the best equilibrium is realized. We will briefly consider three potentially impactful interventions.

6.1 Exaggerating costs

Proposition 1 indicates that, in the model we're considering, there will essentially always be *some* welfare loss due to insufficient social distancing. Even the best equilibrium will involve free-riding, and this alone may motivate a planner to try to push contact levels down by playing up infection costs.

But in some cases the motivation to do so can be extremely acute, because the difference between a scenario with a decent equilibrium and one with only a disastrous "give up on distancing" equilibrium can be razor thin.

Theorem 3. *Consider arbitrary γ, δ, m, and $\underline{c} < 1$. There exists an $\alpha' \in \Re^+$ such that maximal contact is an equilibrium if and only if $\alpha \geq \alpha'$.*

Proof. Theorem 1 describes a disjunction of three conditions, exactly one of which will be operative depending on the values of $\delta(1)$ and γ, and the satisfaction of which will yield $c = 1$ as an equilibrium. Each of the three conditions involves a requirement on α which, after some algebraic manipulation, respectively can expressed as: $\alpha \geq \frac{\ln(\delta(1))(\delta(1)/\gamma)^m}{\ln(\gamma)}$ in the first case, $\alpha \geq 1$ in the second case, and $\alpha \geq \frac{\delta(1)^{\underline{c}m} - \delta(1)^m}{\gamma^{\underline{c}m} - \gamma^m}$ in the third case.

Since all of these place finite lower bounds on α, the equilibrium conditions will be satisfied for all α above some some finite threshold, and none below. □

Figure 9 demonstrates a dramatic case, where a minute decrease in α turns a scenario with a single equilibrium generating worst-possible welfare to one with a single equilibrium that is only moderately suboptimal.

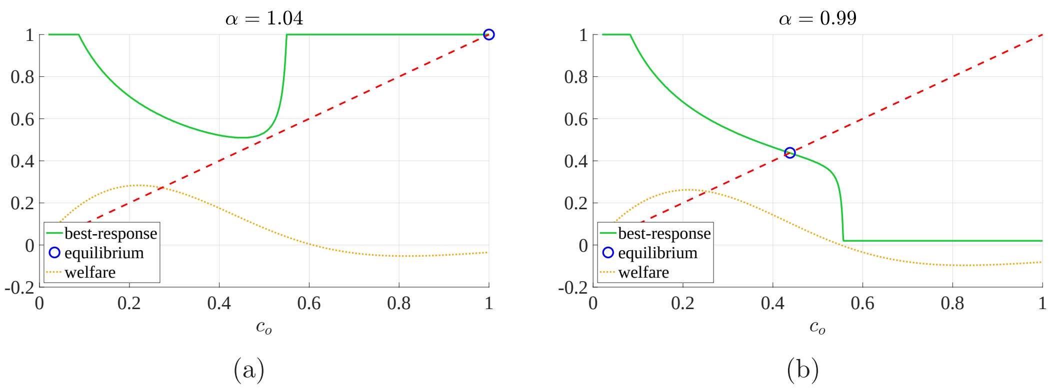

Figure 9: A small decrease in α, the relative weight applied to contact value versus infection risk, transforms a scenario with only a bad equilibrium (a) to one with only a good equilibrium (b). $m = 50$, $\gamma = 0.95$, $\delta(c_o) = 1 - 0.1c_o^{1.2}$.

One can imagine many forms this kind of intervention could take in practice, ranging from outright deception about the ramifications of becoming infected, to actually changing experienced value and costs, say by making indoor dining impossible (diminishing the value of dining-out contacts), or creating a stigma for infection, or even making treatment financially costly.

6.2 Exaggerating the prevailing social distancing rate

Many scenarios will have two or three equilibria rather than just one. In such cases, ending up in "the right" equilibrium can be critical to achieving high welfare. For instance, reconsider Figure 6 — there is a low-contact, high-welfare equilibrium at $c \approx 0.35$ and a maximal-contact, minimal-welfare one ($c = 1$); the third equilibrium has *all* contact levels as a best-response, and is thus unstable and can probably be disregarded. If individuals could be persuaded, either truthfully or otherwise, that "everyone" will be exerting a high-level of restraint and social distancing, perhaps people would converge on the high-welfare equilibrium.

We know from Theorem 2 that the best equilibrium is always the one with lowest contact level. Therefore, the above reasoning may lead to a general tendency to try to bias beliefs towards the notion that prevailing social contact levels are low.

6.3 Enabling lower contact levels

Sometimes a bad maximal-contact equilibrium obtains only because even the minimal feasible contact level is not low enough to avoid the worst costs — imagine a (concave) cost curve, as a function of individual contact level, that is very steep initially and flattens out very quickly; it may only be worth forgoing the value of social contacts if you can avoid that steep cost. In such cases, making lower contact levels feasible may induce a transition from a bad equilibrium scenario to a good equilibrium scenario.

Theorem 4. *Consider arbitrary α, γ, δ, and m with $\delta(1) < \gamma$. There exists a $\underline{c}' \in [0, 1]$ such that maximal contact is an equilibrium if and only if $\underline{c} \geq \underline{c}'$.*

Proof. To prove the theorem it is sufficient to show that an increase in $\underline{c}$ can never remove an equilibrium at $c = 1$. Theorem 1 characterizes the three individually sufficient conditions for equilibrium at $c = 1$. Only the third involves $\underline{c}$ in any way, requiring that $\delta(1) < \gamma$ and $\alpha(\gamma^{\underline{c}m} - \gamma^m) - (\delta(1)^{\underline{c}m} - \delta(1)^m) \geq 0$. Taking the derivative of the left hand side of the latter with respect to $\underline{c}$, we get:

$$m\big(\alpha \log(\gamma)\gamma^{\underline{c}m} - \log(\delta(1))\delta(1)^{\underline{c}m}\big)$$

If $\delta(1) < \gamma$, then this is positive for all values of $\underline{c}$, and thus if $\alpha(\gamma^{\underline{c}'m} - \gamma^m) - (\delta(1)^{\underline{c}'m} - \delta(1)^m) \geq 0$, then $\alpha(\gamma^{\underline{c}m} - \gamma^m) - (\delta(1)^{\underline{c}m} - \delta(1)^m) \geq 0$ for all $\underline{c} > \underline{c}'$, which completes the proof. □

Figure 10 reproduces the example from Figure 3, but in this case also considering an alternative where the minimum number of contacts is 2 rather than 1. In the $\underline{c}m = 2$ case, there is a full-contact equilibrium with devastatingly low welfare, while in the $\underline{c}m = 1$ case there is only a single equilibrium and it is approximately optimal.

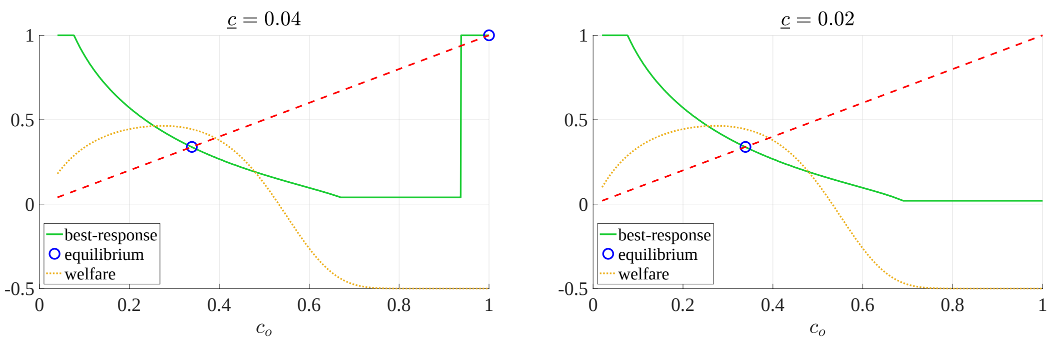

(a) (b)

Figure 10: A small decrease in $\underline{c}$, changing the minimum number of possible contacts from 2 to 1, eliminates the terrible full-contact equilibrium. $m = 50$, $\gamma = 0.8$, $\delta(c_o) = 1 - 0.6c_o^5$

7 Conclusion

Even the most sophisticated attempts to map the course of the COVID-19 pandemic have struggled to predict outcomes with great precision, especially for longer time-horizons [13]. This is perhaps unsurprising when one considers the stubborn uncertainties of viral evolution and government policy-making, and also of behavioral response at the level of people going about their lives.

I have targeted an aspect of that last variable, aiming to hone intuitions about how strategic considerations at the individual level can impact the gross shape of pandemic outcomes: a person may rationally forgo distancing both when others distance very little and when they distance very much (even if not in the middle); relatively small changes in transmissibility or the perceived cost of infection can have dramatically disproportionate effects on behavior; to the extent that people are selfish and well-informed, outcomes will involve *less distancing than is optimal*.

To some readers that last point may seem sufficient to motivate forceful intervention, but I will suggest a couple reasons to be cautious. First, the absence in my model of altruistic motives and behavior, and the absence of positive externality to others from making contact, may present a distorted picture — might individuals even be so concerned about the possibility of infecting others that they *over*-distance in some cases?

More generally, even models that attempt to grapple with every observable high-level feature of this setting will be blind to the many unmeasurable specifics of dilemmas facing individuals and families. Just as models in this domain often fail to predict the actual course of events, well-intended interventions can end up being tragically destructive. When everyone's field of vision is limited, but in different ways, who should set the course?

References

[1] Jude Bayham, Nicolai V Kuminoff, Quentin Gunn, and Eli P Fenichel. Measured voluntary avoidance behaviour during the 2009 a/h1n1 epidemic. *Proceedings of the Royal Society B: Biological Sciences*, 282(1818):20150814, 2015.

[2] Ruggiero Cavallo. Manipulation of would-be free-riders. Working paper, 2019.

[3] Sheryl L. Chang, Mahendra Piraveenan, Philippa Pattison, and Mikhail Prokopenko. Game theoretic modelling of infectious disease dynamics and intervention methods: a review. *Journal of biological dynamics*, 14(1):57–89, 2020.

[4] Frederick Chen. A mathematical analysis of public avoidance behavior during epidemics using game theory. *Journal of theoretical biology*, 302:18–28, 2012.

[5] Maryam Farboodi, Gregor Jarosch, and Robert Shimer. Internal and external effects of social distancing in a pandemics1. *Covid Economics*, page 22, 2020.

[6] Eli P. Fenichel. Economic considerations for social distancing and behavioral based policies during an epidemic. *Journal of health economics*, 32(2):440–451, 2013.

[7] Neil Ferguson, Daniel Laydon, Gemma Nedjati Gilani, Natsuko Imai, Kylie Ainslie, Marc Baguelin, Sangeeta Bhatia, Adhiratha Boonyasiri, ZULMA Cucunuba Perez, Gina Cuomo-Dannenburg, et al. Report 9: Impact of non-pharmaceutical interventions (npis) to reduce covid19 mortality and healthcare demand. `http://www.imperial.ac.uk/media/imperial-college/medicine/sph/ide/gida-fellowships/Imperial-College-COVID19-NPI-modelling-16-03-2020.pdf`.

[8] Michael Greenstone and Vishan Nigam. Does social distancing matter? *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2020-26), 2020.

[9] Donald G. McNeil Jr. How much herd immunity is enough? *The New York Times*, page A6, December 24 2020.

[10] 2020 June 23. White house coronavirus task force members testify on federal response to pandemic. https://www.c-span.org/video/?473229-1/white-house-coronavirus-task-force-members-testify-federal-response-pandemic, 2020.

[11] Kathy Leung, Marcus H.H. Shum, Gabriel M. Leung, Tommy T.Y. Lam, and Joseph T. Wu. Early transmissibility assessment of the n501y mutant strains of SARS-CoV-2 in the United Kingdom, October to November 2020. *Eurosurveillance*, 26(1), 2021.

[12] New and Emerging Respiratory Virus Threats Advisory Group (18 December 2020). Nervtag meeting on SARS-CoV-2 variant under investigation VUI-202012/01. https://khub.net/documents/135939561/338928724/SARS-CoV-2+variant+under+investigation%2C+meeting+minutes.pdf/962e866b-161f-2fd5-1030-32b6ab467896?t=1608470511452, 2020.

[13] Evan L Ray, Nutcha Wattanachit, Jarad Niemi, Abdul Hannan Kanji, Katie House, Estee Y Cramer, Johannes Bracher, Andrew Zheng, Teresa K Yamana, Xinyue Xiong, et al. Ensemble forecasts of coronavirus disease 2019 (COVID-19) in the US. *MedRXiv*, 2020.

[14] Timothy C. Reluga. Game theory of social distancing in response to an epidemic. *PLoS Computational Biology*, 6(5):e1000793, 2010.

[15] Dalmeet Singh Chawla. Critiqued coronavirus simulation gets thumbs up from code-checking efforts. *Nature*, pages 323–324, 2020.

[16] Flavio Toxvaerd. Equilibrium social distancing. Technical report, Faculty of Economics, University of Cambridge, 2020.