NBER WORKING PAPER SERIES

COVID-19, SHELTER-IN PLACE STRATEGIES AND TIPPING

Zhihan Cui
Geoffrey Heal
Howard Kunreuther

Working Paper 27124
http://www.nber.org/papers/w27124

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2020

Support for this research comes from the Alfred P. Sloan Foundation and the Wharton Risk Management and Decision Processes Center. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Covid-19, Shelter-In Place Strategies and Tipping
Zhihan Cui, Geoffrey Heal, and Howard Kunreuther
NBER Working Paper No. 27124
May 2020
JEL No. C72,I12,I18

ABSTRACT

Social distancing via shelter-in-place strategies has emerged as the most effective way to combat Covid-19. In the United States, choices about such policies are made by individual states. Here we show that the policy choice made by one state influences the incentives that other states face to adopt similar policies: they can be viewed as strategic complements in a supermodular game. If they satisfy the condition of uniform strict increasing differences then following Heal and Kunreuther ([6]) we show that if enough states engage in social distancing, they will tip others to do the same and thus shift the Nash equilibrium with respect to the number of states engaging in social distancing.

Zhihan Cui
School of International and Public Affairs
Columbia University
Program on Sustainable Development
New York, NY 10027
USA
zc2322@columbia.edu

Geoffrey Heal
Graduate School of Business
516 Uris Hall
Columbia University
New York, NY 10027-6902
and NBER
gmh1@columbia.edu

Howard Kunreuther
Wharton Risk Management
and Decision Processes Center
The Wharton School
University of Pennsylvania
3819 Chestnut Street, Suite 330
Philadelphia, PA 19104-6366
and NBER
kunreuth@wharton.upenn.edu

1 Introduction

Some governments have responded to the emergence of Covid-19 by undertaking extensive testing and contact-tracing, quarantining those who test positive and their contacts. South Korea, Taiwan and Iceland are in this category. Others are relying on shelter-in -place (s-i-p) orders, which have emerged as one of the most widespread policies for mitigating the spread of Covid-19. (For a review of Covid-19-related policies see Dalton et al. [3].) In general these policies are implemented at the federal or equivalent level. In this regard, the national governments of the UK, Italy, France and South Korea have all mandated[1] that most of their residents should remain at home.[2] Rather uniquely, the U.S. has left state governors to choose whether to implement such policies. As a result the majority of states, but not all of them, now have such orders in place, and the choice has become a political one, with most Democratic governors implementing such orders but many - though not all - Republican governors reluctant to do so.

Testing plus contact tracing and s-i-p orders can both be seen as different implementations of social distancing. S-i-p orders are a broadly targeted form of social distancing, requiring individuals or households to be physically separate from others, whereas testing plus contact tracing and isolation are a much more targeted form of social distancing. The ultimate goal of both approaches is still to prevent Covid-19 from spreading by separating people who might be carriers of the virus from the rest of the population..

S-i-p orders have costs and benefits (see Thunstrom et al. for a cost-benefit analysis [9]). The costs are obvious and largely economic: they bring the local economy to a grinding halt, as many businesses cannot continue to operate in a world of s-i-p orders. There are also social costs associated with isolation and lack of social interactions. There are health benefits since

[1]In South Korea this is actually a government recommendation, which is very widely followed, rather than an order. Thanks to Jisung Park for this point.

[2]And in some cases wear masks in public places. On wearing masks, see Sunstein [8].

illnesses spread much less rapidly and fatalities are reduced when most people are required to stay at home. With no social distancing regulations in place, the average person with Covid-19 in the United States will infect about 3 others, whereas with the aggressive social distancing practiced in New York State, that number appears to have fallen to under one.[3] (For a discussion for the data for New York City see Harris [5], and for a general discussion of social distancing in epidemic models see Kelso et al.[7], who analyze how social distancing can reduce the rate at which a disease spreads from infected to susceptible populations.)

This note shows formally that a state's decision on whether to introduce shelter-in -place regulations in the U.S. depends on how many other states have already instituted such orders. The larger the number of states with s-i-p policies, the more effective a new one is and more likely it is that a new state will follow suit. More formally, state $i's$ payoff from implementing an s-i-p order depends on the choices of states j for the following reason: if state j does **not** implement such a policy, then the virus can continue to spread in state j and people who travel between j and i can infect people in i, undercutting $i's$ shelter-in -place policy.

A good illustration is provided by the tri-state area of New York, New Jersey and Connecticut. Residents of all these states commute to and work in New York City, meaning that if New York closes down its businesses, residents of all three states are affected. Many residents of New Jersey and Connecticut will lose their jobs and thus have less reason to travel to New York, and businesses in those states will close too. So a move by New York to have people shelter-in -place and to close businesses will make it easier for the governors of adjacent states to do likewise: the incremental economic costs are lower because part of the work was already done by New York. The fact that so many people in the tri-state area travel between states for work, shopping and entertainment, also illustrates well the ease with which a virus

[3] According to Governor Andrew Cuomo at briefing on 4/19/2020 it is currently 0.9.

can spread from one state to another. Reducing the incidence of a diseases in one state will reduce its incidence in others with whom residents of the first state interact.

Given that the spreading of a virus depends not only on a state's own action but those of others, the decisions on whether or not to implement s-i-p orders by individual states can be formalized as a game. This particular game is supermodular and so will have multiple Nash equilibria, including a greatest and a least equilibrium (Topkis [2]). If the effectiveness of an s-i-p policy in state i depends on whether such orders are in place elsewhere and increases with this number, then the game between states is characterized by social reinforcement, and in particular its payoffs may show what Heal and Kunreuther ([6]) call uniform strict increasing differences, a strong form of strategic complementarity.

The next section models this interdependence and shows how the existence of *tipping sets* arises. A tipping set is a set of players with the following property: if all member of this set choose to implement shelter-in -place policies, then the best response of every other agent will be to follow suit. So the member of the tipping set can drive all others to the adoption of shelter-in -place policies, even in the absence of a federal mandate for such policies.

One can also have local tipping sets. In the context of the social distancing problem facing states, the Nash equilibria may be regional rather than national, so that if one or more states change their strategy, some nearby states may follow suit. For example, a change in policy by New York may force New Jersey and Connecticut to do likewise. Similarly there may be strong links between Georgia, South Carolina and Tennessee. Proximity does not necessarily have to be geographic: it could be measured in terms of economic links between the states.

The idea that one state's policies reinforce those of another can be tested empirically, though we are not aware of any completed studies on this question. There are published data on the incidence of Covid-19 by state and date and also on when various s-i-p measures in each state or municipality are introduced.[4] We are now using these data determine whether there are positive effects of one state's policies on others.

There are other more complex elements of the relationships between states. They do assist each other in attaining health goals through the reinforcement we have discussed, but they also compete for scarce medical equipment such as personal protective equipment and ventilators, bidding up prices. New York Governor Andrew Cuomo has frequently complained in his press briefings of the lack of a centralized national purchasing policy and the way in which this pits states against each other. This means that one state's actions in response to Covid-19 raises the costs of the actions that others wish to take. This behavior is a result of policies for obtaining medical equipment to deal with illnesses from Covid-19 and not due to social distancing policies.

As mentioned above, there is also a political dimension to choices in dealing with the coronavirus pandemic. Democratic governors more likely than their Republican counterparts to recognize the seriousness of Covid-19 and the need for collective action to mitigate the spread of the virus. This could be captured in differences in agents' preferences as we demonstrate in the formal game theoretic model presented in the next section.

Heal and Kunreuther ([6]) provide a simple example of a game that meets all the conditions mentioned above. There are I players and each may choose as a strategy either zero or one: think of zero as no policy and one as an s-i-p policy. The payoff to choosing zero, is always 0.1. The payoff to agent j of choosing 1 is equal to the number of others who choose 1. If no one else chooses 1, the payoff is 0. It then increases linearly depending on how many others choose 1 so if n agents choose 1, the payoff to the $n+1-th$ agent to making this choice is n. In this particular game there are only two Nash equilibria: every agent chooses 0 or every agent chooses 1. If every agent has

[4]From Kinsahealthcare and the CDC's data on influenza-like illnesses.

chosen 0 and a single agent switches to 1, then all the other agents will also want to switch to 1. In other words, the game has been "tipped" from a Nash equilibrium where everyone chooses 0 to a Nash equilibrium where everyone chooses 1.

2 Formal Model

There are I agents (states) indexed by $i = 1, 2, ..., I$. Each has a strategy s_i and a strategy space given by two alternatives $S = \{0, 1\}$ where $s_i = 0$ denotes no s-i-p or social distancing policy and $s_i = 1$ indicates that such a policy is in place. The vector $S \in R^I$ represents the list of strategies chosen by all agents $S = (s_1, s_2, ...s_I)$. Each agent's payoff function $U_i(S) : S^I \to R^1$ depends on the choices of all agents, its own and those of others. We let 0_i or 1_i denote a zero or a one in the $i-th$ position of S and the vector S_{-i} be the vector of all choices made by states other than i. We assume that the U_i all satisfy uniform strict increasing differences, that is using the usual vector ordering on R^I, $\exists \epsilon > 0 : \ S'_{-i} > S_{-i} \ \Rightarrow$

$$U_i\left(1_i, S'_{-i}\right) - U_i\left(0_i, S'_{-i}\right) \geq \epsilon + U_i\left(1_i, S_{-i}\right) - U_i\left(0_i, S_{-i}\right) \tag{2.1}$$

In words, consider two configurations of strategy choices by players other than i, denoted S_{-i} and S'_{-i}. Then if in S'_{-i} at least one state has changed from zero to one relative to S_{-i}, which is implied by $S'_{-i} > S_{-i}$, then the payoff to state i to changing from zero to one is strictly and uniformly greater at S'_{-i} than at S_{-i}. This means that agent j changing from zero to one raises the payoff to this change for agent i for any i and j. This is implied by the interactions between state strategies discussed above: the adoption of an s-i-p policy by state j makes such a policy more effective for state i. In the inequality (2.1) the parameter ϵ is a measure of the degree of social reinforcement: the greater is ϵ, the greater is the degree of social reinforcement or strategic complementarity and the small is the tipping set.

For simplicity we are assuming the ϵ to be independent of the states involved, though the discussion above of the tri-state area makes it clear that in reality some pairs of states reinforce each other more than other pairs. Think of New York and New Jersey versus New York and Alabama.

Tipping sets are important in this analysis. Intuitively a tipping set is a subset T of players which has the following property. If all the members of T choose strategy 1, then the best response for any other player is strategy 1. If all members of T choose s-i-p orders, then every other state finds that its best strategy is also to choose an s-i-p order. Formally, if $S_i = 1 \forall i \in T$, then $\forall i \notin T,\ U_i(1_i, S_{-i}) \geq U_i(0_i, S_{-i})$. A minimal tipping set is a tipping set with the property that no strict subset is also a tipping set.

The set of possible strategy vectors S in this game is the set of vectors of the form $(0, 1, 1, 0, 0, ...)$ where every coordinate is a zero or a one. These vectors form the vertices of the unit cube in R^I, which is a lattice. By assumption (2.1), the game is supermodular. Hence we know by a theorem of Topkis ([2]) that the set of pure strategy Nash equilibria is non-empty and contains greatest and least elements which we call $\bar{S}$ and $\underline{S}$ respectively. From Dall, Lakshmivarahan and Verma ([4]) we know that for two players $\bar{S} = (1, 1)$ and $\underline{S} = (0, 0)$ (corollary 3.2) and for three players these are $(1, 1, 1)$ and $(0, 0, 0)$ (Corollary 3.6). For two and three players, then, the greatest and least Nash equilibria are where all agents choose 1 or all choose 0. We assume this is also true for I players: the maximal Nash equilibrium is where all players choose 1 and the minimal where they all choose 0. In the Appendix we will give simple conditions that are necessary and sufficient for this to be the case.

Under these conditions, we can prove that there is a tipping set T of states with the ability to tip the no-s-i-p order equilibrium to the all-s-i-p equilibrium. Furthermore there is a tipping set that will tip *any* equilibrium with less than every state having s-i-p orders to one where all do so. It is also true that our proof that there is a set that will tip the equilibrium of all zeros to that of all ones applies with minor modifications to showing that there is a set that will tip from the least Nash equilibrium to the greatest, whatever these may be. A formal statement of our results is given in the technical appendix, together with proofs.

In addition to tipping, we can have the related phenomenon of cascades. A cascade occurs when a change of policy by agent 1 causes 2 to change her policy, which in turn causes 3 to change and so on, a classical "domino effect." This process may take in all agents or only a subset. A simple example from Heal and Kunreuther ([6]) is as follows. There are 10 agents. For any agent the return to setting $s_i = 0$ is $0.91i$. The return to $s_i = 1$ is $\#(1)$, the number of other agents also choosing one. Clearly all zeros and all ones are both Nash equilibria. Suppose that all are choosing zero and agent 10 decides to switch to one. Then the return to agent 1 to choosing 1 is now $1 > 0.91i$ and she will switch to 1. Agent 2 will now find that the return to choosing 1 is $2 > 1.82$ and will switch. And so on for all agents up to and including 9. Agent 10, by switching, started a cascade of all the other agents beginning with 1. Heal and Kunreuther ([6]) give sufficient conditions for a cascade to occur. It is possible that the connections between New York and adjacent states are best described by a cascade rather than by tipping.

Returning to the issue of political differences on s-i-p policies, it is possible we could model these by differences in the states' payoff functions $U_i(S_{-i}, S_i)$: republican states may value the outcomes associated with s-i-p policies - reduced morbidity and mortality - but have a preference against the action of implementing an s-i-p policy. The might prefer a world in which good public health outcomes are attained by other states implementing s-i-p policies while they don't: they strongly prefer $(1_{-i}, 0_i)$ (the vector of ones everywhere except in the $i-th$ position to $(1, 1, ..., 1)$, the vector of all ones. In this case there can be no Nash equilibrium where all agents choose one: the greatest Nash equilibrium $\bar{S}$ will satisfy $\bar{S} < (1, 1, ..., 1)$. The fact that states with conservative governors, such as Georgia, are moving first to relax s-i-p policies, is consistent with them having a strong negative preference for these policies. The importance of political orientation for attitudes towards Covid-19 is studied by Barrios and Hochberg ([1]), who show that the attention paid to Covid-19 is negatively correlated with support for Donald Trump in the last presidential election. Using Google search data, they show that areas showing high Trump support only start paying attention when there are Covid-related deaths in their region, or when prominent conservative figures emphasize the reality of the epidemic. Their work actually suggests that there is support for social distancing in conservative states, but only once lives are being lost. They suggest that preferences evolve over the course of the epidemic, not something we can model in our framework.

3 Conclusions

shelter-in -place (s-i-p) strategies and social distancing are integral to overcoming a pandemic. In the U.S. s-i-p strategies have to be implemented by states, which face complex combinations of costs and benefits from their possible choices. Their decisions are affected by those of other states since strategy choices demonstrate social reinforcement. A compelling illustration of this interdependence is the interactions between New York and its neighboring states: the tri-state region can be seen as a single unit in terms of employment, commuting, entertainment and retail shopping. A move towards social distancing by any of these states will affect the other two, and its effectiveness will depend on the reactions of the others. Because of this, we can model their choices as a game. Specifically, we show that the choice of an s-i-p policy by a single state can tip a system to a new Nash equilibrium at which many more agents have adopted shelter-in -place or social distancing policies. It could also cause a cascade from one equilibrium to another.

4 Appendix

Theorem 1. *Under assumption (2.1), there is a minimal tipping set T consisting of less than $I-1$ agents, which will tip the least Nash equilibrium to the greatest Nash equilibrium. Furthermore, any Nash equilibrium with less than $I-1$ s-i-p orders or social distancing choices can be tipped to the equilibrium with I such orders.*

Proof. We study the effect on agent $j's$ payoff of changing from no s-i-p to an s-i-p (changing from 0 to 1) and how this effect is altered by changes in the strategy choices of another agent i. We know by (2.1) that if i switches from 0 to 1 then this will increase the incremental payoff to j from the same switch. Let $S_{-i-j}, 1_i, 0_j$ denote the vector of strategies in which all agents other than i, j are choosing $S_k \in S_{-i-j}$ and i, j are choosing 1 and 0 respectively. (S_{-i-j} is the vector of strategies chosen by all agents other than i and j.) Define

$$\Delta_j\left(i=0, S_{-i-j}\right) = U_j\left(S_{-i-j}, 0_i, 1_j\right) - U_j\left(S_{-i-j}, 0_i, 0_j\right) \tag{4.1}$$

and

$$\Delta_j\left(i=1, S_{-i-j}\right) = U_j\left(S_{-i-j}, 1_i, 1_j\right) - U_j\left(S_{-i-j}, 1_i, 0_j\right) \tag{4.2}$$

These are the returns to j from changing from 0 to 1 when i chooses either 0 (first line) or 1 (second line) and everyone else chooses $s_k \in S_{-i-j}$. The difference between these is

$$\Delta_{ij}\left(S_{_{-i-j}}\right) = \Delta_j\left(i=1, S_{-i-j}\right) - \Delta_j\left(i=0, S_{-i-j}\right) \geq 0 \tag{4.3}$$

This is the increase in the return to $j's$ change of strategy as a result of $i's$ change of strategy and from (2.1) we know that this is positive. We focus on equation (4.3) when all agents other than i and j are choosing strategy 0 so as to derive conditions for tipping the Nash equilibrium of all zeros to that of all ones:

$$\Delta_{ij}(0) = \left\{U_j\left(0^{I-2}, 1_i, 1_j\right) - U_j\left(0^{I-2}, 1_i, 0_j\right)\right\} - \left\{U_j\left(0^{I-2}, 0_i, 1_j\right) - U_j\left(0^{I-2}, 0_i, 0_j\right)\right\} \tag{4.4}$$

where 0^{I-2} indicates that there are $I-2$ zeros in position other than i and j. Consider the following sequence of inequalities, which link the equilibrium with all 0s to that will all 1s in a series of steps in each of which an additional state changes strategy from 0 to 1, and which hold because of (2.1):

$$U_i\left(0^{I-1}, 1_i\right) - U_i\left(0^{I-1}, 0_i\right) + \epsilon < U_i\left(0^{I-2}, 1_1, 1_i\right) - U_i\left(0^{I-2}, 1_1, 0_i\right) \tag{4.5}$$

$$U_i\left(0^{I-2}, 1_1, 1_i\right) - U_i\left(0^{I-2}, 1_1, 0_i\right) + \epsilon < U_i\left(0^{I-3}, 1_1, 1_2, 1_i\right) - U_i\left(0^{I-3}, 1_1, 1_2, 0_i\right)$$

$$U_i\left(1_1, ..1_{I-2}, 0_j, 1_i\right) - U_i\left(1_1, ..1_{I-2}, 0_j, 0_i\right) + \epsilon < U_i\left(1_1, ..1_{I-1}, 1_i\right) - U_i\left(1_1, ..1_{I-1}, 0_i\right)$$

The first inequality here (4.5) shows that the payoff to state i from a strategy change is raised by at least ϵ when state 1 also picks strategy 1. The second inequality shows that the payoff to i from the change is again increased by ϵ when state 2 also changes from 0 to 1. Working back from a typical inequality in this sequence we find that

$$U_i\left(0^{I-k}, 1_1, 1_2, ..., 1_i\right) - U_i\left(0^{I-k}, 1_1, 1_2, ..., 0_i\right) > (k-1)\epsilon + U_i\left(0^{I-1}, 1_i\right) - U_i\left(0^{I-1}, 0_i\right)$$

Note that $U_i\left(0^{I-1}, 1_i\right) - U_i\left(0^{I-1}, 0_i\right) < 0$ as the vector of zeros is a Nash equilibrium so zero is a best response. Note also that the last difference in this sequence $U_i\left(1_1, 1_2, ...1_{I-1}, 1_i\right) - U_i\left(1_1, 1_2, ...1_{I-1}, 0_i\right) > 0$ as the vector of all ones is a Nash equilibrium and therefore 1 is a best response. As the sequence of differences starts negative and ends positive it must change sign: there will be a $k < I-1$ such that $(k-1)\epsilon - U_i\left(0^{I-1}, 1_i\right) + U_i\left(0^{I-1}, 0_i\right) > 0$ and the first k states will form a tipping set. To be precise we need k to satisfy

$$(k-1)\epsilon > U_i\left(0^{I-1}, 1_i\right) - U_i\left(0^{I-1}, 0_i\right) \ \forall i \tag{4.6}$$

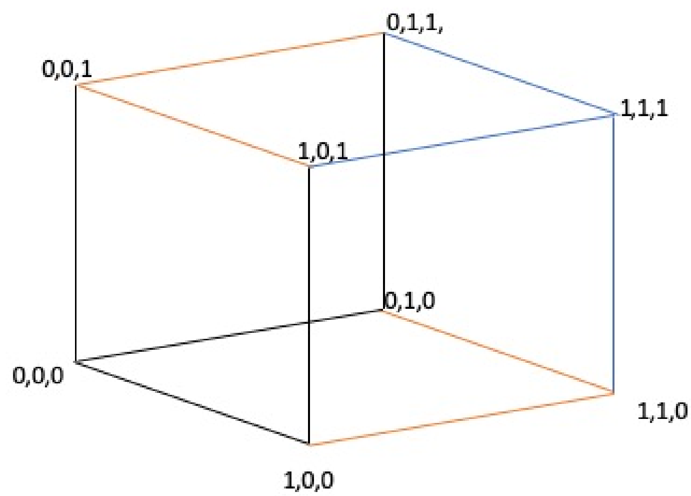

Figure 4.1: All possible plays for three players. Black edges are connected to the least Nash equilibrium and blue to the greatest.

In this case each of the other states finds it in its interest to change its strategy from zero to one and the equilibrium of zeros is tipped to that of ones if the first k states all change from zero to one. Equation (4.6) shows a tradeoff between the social reinforcement parameter ϵ and the size of a tipping set k : the great the social reinforcement (the greater ϵ) the smaller the number k in the tipping set. $\square$

Next we turn to the characterization of the greatest and least Nash equilibria of the game $\bar{S}$ and $\underline{S}$, whose existence is assured by the theorem of Topkis ([2]).

Theorem 2. *A necessary and sufficient condition for* $\underline{S} = (0, 0, ..., 0)$ *and*

$\bar{S} = (1, 1, ..., 1)$ *is that for every agent i, if all other agents have chosen the same strategy s, then that common strategy s is i's best response.*

Proof. The proposition is immediate. □

In plain English, we have Nash equilibria at all zeros and all ones if it never pays to be the odd-man-out. Proposition 2 has implications in terms of the structure of agents' utility functions. It requires that $U_i(0_i, 1_i) - U_i(0_i, 0_i) < 0$ and $U_i(1_i, 1_i) - U_i(1_i, 0_i) > 0$. So the derivative of $i's$ payoff with respect to its strategy depends heavily on the strategy choices of others, to the extent of changing sign if these other strategy choices all change.

Figure 4.1 illustrates how payoffs to a choice by i vary with the choices of others. There are three players and $(0, 0, 0)$ and $(1, 1, 1)$ are Nash equilibria. So if starting from $(0, 0, 0)$ agent 1 changes to 1 and moves to vertex $(1, 0, 0)$ then she is worse off. Likewise if agent 2 moves to $(0, 1, 0)$ she is worse off. However if agent 1 changes from zero to one starting from $(0, 0, 1)$ and so moves from vertex $(0, 0, 1)$ to vertex $(1, 0, 1)$ she may gain.

References

[1] John Barrios and Yael Hochberg. Risk perception through the lens of politics in the time of the covid-19 pandemic. Working Paper 27008, NBER, April 2020.

[2] Topkis D. Equilibrium points in nonzero-sum n-person sumodular games. *SIAM Journal of Control and Optimization*, 17:733–787, 1979.

[3] Craig Dalton, Stephen Corbett, and Anthea Katelaris. Pre-emptive low cost social distancing and enhanced hygiene implemented before local covid-19 transmission could decrease the number and severity of cases. Technical report, Social Science Research Network, Available at SSRN: https://ssrn.com/abstract=3549276 or http://dx.doi.org/10.2139/ssrn.3549276, March 2020.

[4] Sudarshan Dhall, S Lakshmivarahan, and Pramode Verma. On the number and distribution of nash equilibria in supermodular games and their impact on the tipping set. In *DOI: 10.1109/GAMENETS.2009.5137462. Source: IEEE Xplore*. 2009.

[5] Jeffrey E. Harris. The corona virus epidemic curve is already flattening in New York City. Working Paper 26917, NBER, April 2020.

[6] Geoffrey Heal and Howard Kunreuther. Social reinforcement: Cascades, entrapment, and tipping. *American Economic Journals: Microeconomics*, 2(1):86–99, 2010.

[7] Joel Kelso, George Milne, and Heath Kelly. Simulation suggests that rapid activation of social distancing can arrest epidemic development due to a novel strain of influenza. *BMC Public Health 9, 117 (2009).*, 9(117), 2009.

[8] Cass Sunstein. The meaning of masks. *Journal of Behavioral Economics and Policy*, Forthcoming, 2020.

[9] Linda Thunstrom, Stephen Newbold, David Finnoff, Madison Ashworth, and Jason F. Shogren. The benefits and costs of using social distanc-ing to flatten the curve for covid-19. *Journal of Benefit-Cost Anal-ysis.*, page Available at SSRN: https://ssrn.com/abstract=3561934 or http://dx.doi.org/10.2139/ssrn.3561934, Forthcoming.