The Behavioral SIR Model, with Applications to the Swine Flu and COVID-19 Pandemics

Samuel Engle Jussi Keppo Marianna Kudlyak*
Elena Quercioli Lones Smith Andrea Wilson

NY Fed Seminar, 4PM on 6/30/2020

(*Views expressed are not those of the Federal Reserve Bank of San Francisco or the Federal Reserve System.)

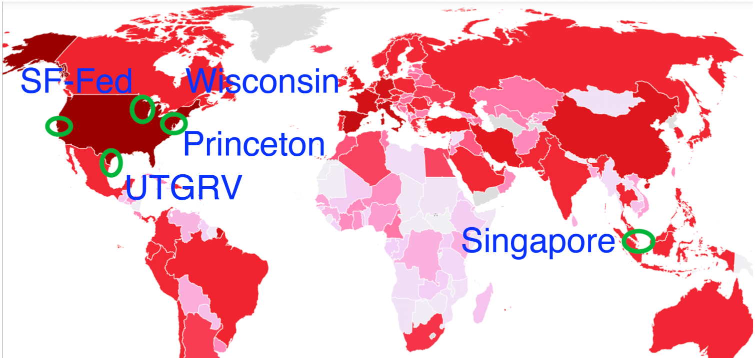

Plan of Talk: Deriving and Using a New Contagion Model

- The influential SIR contagion model
  1. is linear, and so tractable
  2. makes extreme predictions, especially later on in a contagion
  3. ignores human behavior
- We create a Behavioral SIR (BSIR) model that
  1. accounts for optimal avoidance behavior in a Nash equilibrium
  2. is log-linear, and so still tractable for analysis
  3. makes less extreme predictions (consistent with COVID so far)
  4. subsumes the SIR model as a special case for low infectiousness or small disease losses (crucial for statistical tests)
- For COVID19 and Swine Flu (2009), we reject the SIR model
- For COVID19, our BSIR model make sense of time series properties in countries and states, pre- and post-lockdown
- Data from the Swine Flu allows us to evaluate the BSIR model through the entire course of the contagion to herd immunity

- Contagion math in the best of times depends on
  1. Biology: how infectious is the infection?
  2. Sociology: networks, segregation, "Super spreaders"
  3. Geography: meeting rates are higher in dense cities
  4. Culture: in Italy, the kiss sometimes replaces the handshake
  5. Game theory: how we react to payoffs and each other
  6. Political economy: Do lockdowns or stay-in-place work? Are people responsive?

SI / SIS / SIR

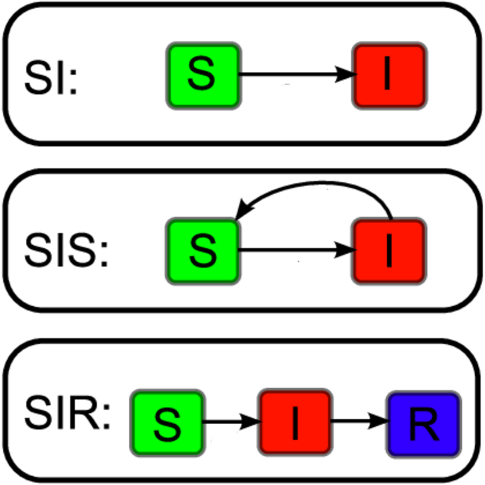

The SIR Model (1927)

- The model takes place in continuous time $t \in [0, \infty)$
- Population is the continuum $[0, 1]$ (no aggregate randomness)
- State transition process of people in the SIR model
- Mass $\sigma(t)$ of individuals are *susceptible* to a disease
- *prevalence* $\pi_t \in (0, 1)$ is the mass of contagious individuals
  - Given: seed mass $\pi_0 > 0$, with $\sigma_0 = 1 - \pi_0$
- *Incidence* is the inflow of new infections
- The *passing rate* is the mean number $\beta > 0$ of susceptible people per unit time each contagious person infects
  - β increases in disease contagiousness, population density
  - β reflects culture and social networks.

The SIR Model

- Anyone infected gets better (or dies) at *recovery rate* $r > 0$.
- *random and independent meetings* $\Rightarrow$ incidence is $\beta\sigma\pi$

$$\begin{aligned}\dot{\sigma}(t) &= -\text{incidence} = -\beta\pi(t)\sigma(t) \\ \dot{\pi}(t) &= \text{incidence} - \text{recoveries} = \beta\pi(t)\sigma(t) - r\pi(t)\end{aligned}$$

Lemma

The susceptible mass $\sigma(t)$ monotonically falls, and prevalence $\pi(t)$ first rises and then falls.

- Proof: $\dot{\pi}(t) = [\beta\sigma(t) - r]\pi(t)$
- A mass ρ is *recovered/removed* and immune
- We ignore $\rho(t)$, as it does not impact dynamics: $\dot{\rho}(t) = r\pi(t)$

Herd Immunity

- *Herd immunity*: Epidemic dies out when enough of the population is immune (high ρ) that its spread stops naturally because too few people can transmit it (low σ)

- tipping point $\Leftrightarrow \dot{\pi}(t) \leq 0 \Leftrightarrow \beta\hat{\sigma}\hat{\pi} = r\hat{\pi}$.

$\Rightarrow$ *basic reproduction number* **R0** $\equiv \beta/r$.

Lemma

Herd immunity happens if $\beta\sigma\pi \leq r\pi \Leftrightarrow \sigma \cdot R0 \leq 1$.

- Published COVID estimates $R0 = 2.3 \Rightarrow \rho_t > 1 - 1/2.3 \approx 0.56$

- "Newsom projection: 56% of California would be infected in 8 weeks without mitigation effort" (2020/03/19)

Goal: Marry Economics and Epidemiology

Our Goal: Behavioral SIR Model

- There is some $\underline{\pi} > 0$ and $0 < \varphi < 1$ such that:
  - If $\pi \leq \underline{\pi}$, SIR dynamics apply (our "chill" regime)

$$\begin{aligned} \dot{\sigma}(t) &= -\beta q(\pi)\sigma(t)\pi(t) \\ \dot{\pi}(t) &= \beta q(\pi)\sigma(t)\pi(t) - r\pi(t) \end{aligned}$$

  - If $\pi > \underline{\pi}$, then the "vigilant" regime obtains:

$$\begin{aligned} \dot{\sigma}(t) &= -\beta q(\pi)\sigma(t)\underline{\pi}^{1-\varphi}\pi(t)^{\varphi} \\ \dot{\pi}(t) &= \beta q(\pi)\sigma(t)\underline{\pi}^{1-\varphi}\pi(t)^{\varphi} - r\pi(t) \end{aligned}$$

- $\underline{\pi} \geq 1$, then only the SIR dynamics obtain.

Incentives Matter in Contagions

- A disease does not pass the same
  1. among humans or animals in the SIR model.
  2. among chill people as alert
     - Example: Measles outbreaks have much higher infection rate than measles pandemics.
- We will focus on optimizing strategic behavior, since it can change very rapidly in the contagion

Incentives Matter in Contagions

- A *huge and longstanding literature* in epidemiology (including some economists lately!) posits exogenous ways that people modify reduce the passing rate as the contagion worsens.
- This is like the adaptive expectations literature of the 1960s.
- The Lucas Critique: must close the loop with equilibrium
  - disease prevalence rises $\Rightarrow$ more vigilant $\Rightarrow$ realize others are more vigilant $\Rightarrow$ relax (strategic substitutes)
  - equilibrium fully accounts for this (infinite) feedback cycle.
  - no arbitrary adjustment rule works
- We build on the model of "Contagious Matching Games" (2006 Quercioli and Smith), where people best reply to a prevalence, which acts like a price in an "implicit market"

Passing Games

- Counterfeit money vs disease: unwitting sharing of a rival "bad" vs unwitting sharing nonrival "bad"

- We build on the model of "Contagious Matching Games" (2006 Quercioli and Smith), where people best reply to a prevalence, which acts like a price in an "implicit market"

The Contagious Matching Game

- World's biggest game: Everybody is a player $[0, 1]$
- The highest stake game: life of death (or sickness): loss L
- Action: Vigilance $v \geq 0$ costs v and reduces the passing rate
- Players minimize expected total losses

Some Motivation for Our Model

Most say people's actions affect spread of COVID-19

% who say the actions of ordinary Americans affect how the coronavirus spreads in the U.S. ...

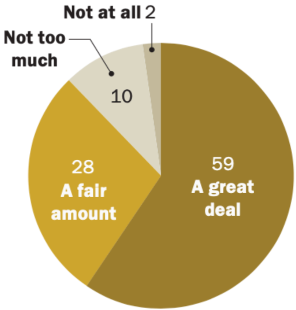

Note: No answer responses not shown.
Source: Survey of U.S. adults conducted June 16-22, 2020.

PEW RESEARCH CENTER

Some Motivation for Our Model

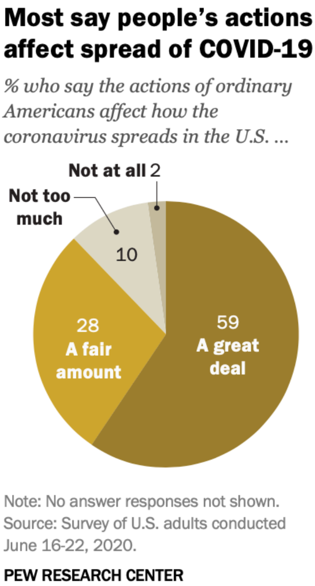

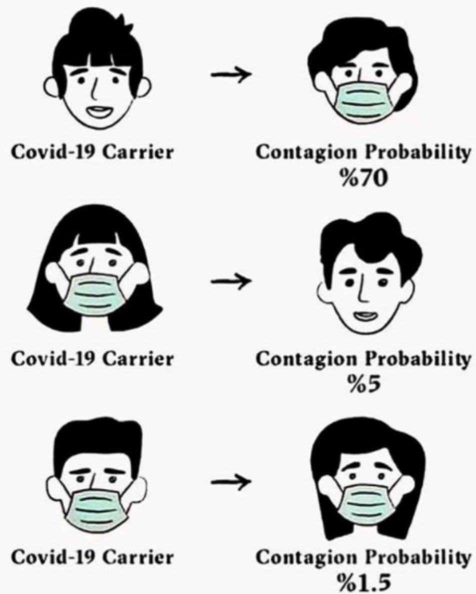

How Vigilance Reduces Passing: the Filter function

- *Filter function* $f(v) \in [0,1]$ linearly scales down passing rates

$\Rightarrow$ Passage rate is $\beta f(v) f(w)$ if vigilance v contagious person just meets vigilance w susceptible people

$\Rightarrow$ diminishing returns: $f(0)=1>0=f(\infty)$ & $f' < 0 < f''$.

- A symmetric function is a simplifying assumption
  - Intensive margin: a mask is equally protective of both parties.
  - Extensive margin: Not meeting also symmetrically protects both parties — $f(v)$ = fraction of meetings one keeps
- This multiplicative (log-modular) form is for simplicity.
- A vaccination is easy vigilance: one jab $\Rightarrow$ nearly perfect filter
- Posit hyperbolic filter function $\boxed{f(v) = (1+v)^{-\gamma}}$, for $\gamma > 0$
  - γ = *filter elasticity* in terms of "total vigilance" $V = 1 + v$.
  - 1% more total vigilance leads to γ% infection risk reduction

Vigilance Optimization

- People are first obliviously contagious, and next knowingly so.
- $\pi =$ mass of unaware contagious individuals
- A *potentially susceptible* if infected with chance $q(\pi) = \frac{\sigma}{\sigma+\pi}$.
- Potentially susceptible people minimize selfish expected total losses:

$$\beta f(v) E[f(W)] q(\pi) \pi L + v$$

- $f' < 0 < f'' \Rightarrow \exists$ a corner solution or a unique interior optimum.
- Since everyone makes the same choice, only pure strategy symmetric Nash equilibria exist, with $W = v^* \geq 0$.
- *flow disease loss* as $\delta(v, v^*, \pi) = \beta f(v) f(v^*) q(\pi) \pi L$

Individual Optimality in Equilibrium

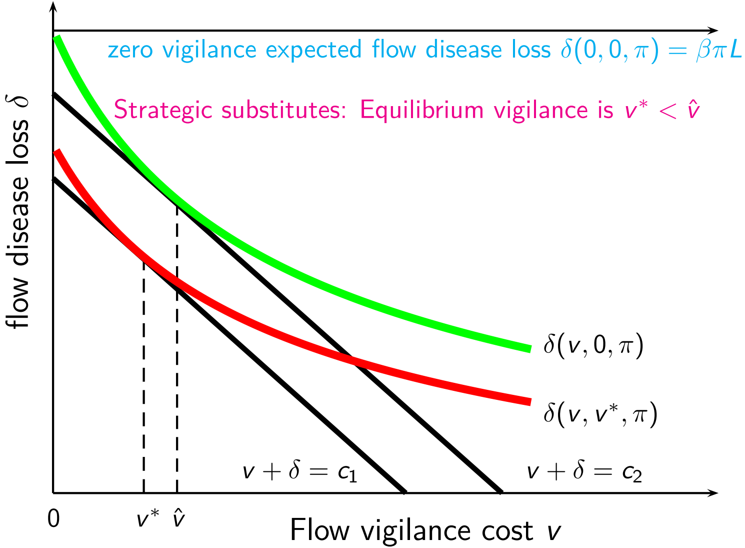

Nash Equilibrium Vigilance

- Vigilance vanishes for low prevalence $\pi \leq \underline{\pi}$, where

$$\underline{\pi} \approx [\beta L(1-\varphi)/(2\varphi)]^{-1}.$$

where $\varphi \equiv 1/(2\gamma+1)$ *does not depend on L, β*

Theorem

There is a unique Nash equilibrium for any prevalence $\pi \geq 0$. Equilibrium vigilance $v^*(\pi)$ vanishes for $\pi \in [0, \underline{\pi}]$, and is increasing for $\pi \geq \underline{\pi}$, for some prevalence threshold $\underline{\pi} > 0$ that is falling in L and β, but rising in φ.

- Note: Any dynamic equilibrium of a continuum agent game requires static Nash play every period

$\Rightarrow$ The only assumption here is a constant loss L, which holds if
  - people are motivated by current losses, or
  - people are forward-looking but act as if in a steady-state.
  - Dynamics impossibly hard to forecast — even experts disagree

Nash Equilibrium Passing Rate

- The *behavioral passing rate* $B(\pi|\varphi) = \beta f(v^*)^2$ is the innate passing rate β times any two individuals' equilibrium filter.

Theorem

The behavioral passing rate has two regimes:

$$B(\pi|\varphi) = \begin{cases} \beta & \pi \leq \underline{\pi} \quad (\textit{chill}) \\ q(\pi)\beta(\underline{\pi}/\pi)^{1-\varphi} \approx \beta(\underline{\pi}/\pi)^{1-\varphi} & \pi > \underline{\pi} \quad (\textit{vigilant}) \end{cases}$$

Given our filter, we have

$$\text{incidence rate} = \mathcal{SI} \text{ meeting rate} \times \text{ passing chance}$$

$$\begin{aligned} \Rightarrow \textit{incidence-prevalence elasticity} &= 1 + \text{passing rate elasticity} \\ &= 1 + (\varphi - 1) \\ &= \varphi \end{aligned}$$

Incidence-Prevalence Elasticity

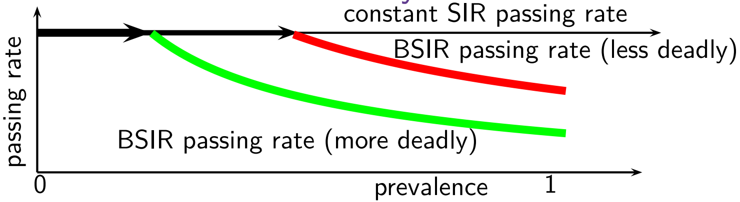

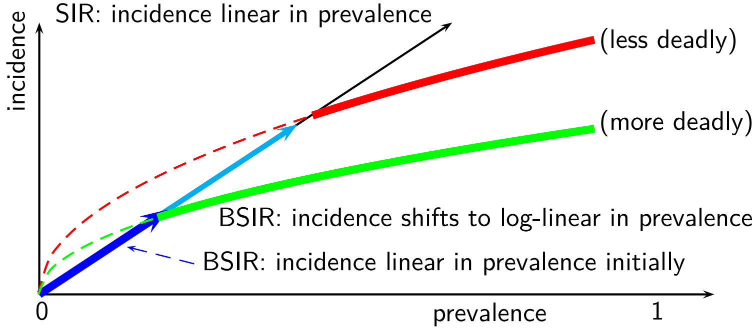

Incidence-Prevalence Elasticity

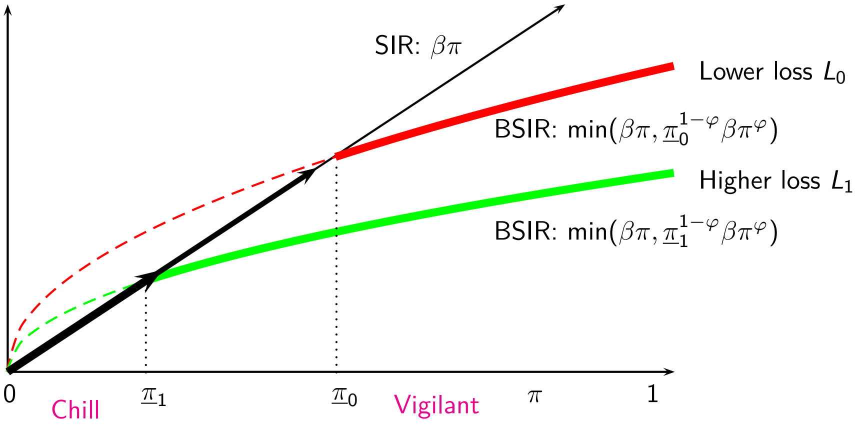

Prevalence Elasticity of Incidence

Corollary (Breakout Incidence)

Equilibrium incidence $B(\pi)\pi\sigma$ is log-linear in prevalence $\pi \geq \underline{\pi}$,

$$\log(\textit{incidence}) = \log[B(\pi)\pi\sigma] = b + \varphi \log \pi + \log \sigma$$

where the incidence-prevalence elasticity is $\varphi \equiv 1/(2\gamma + 1) < 1$, and the intercept b increases in φ and β, and falls in L.

Corollary

For the same number of cases, the passing rate rises in population.

Assume a lockdowns reduce β, by foreclosing on opportunities.

Corollary

Lockdown parallelly shifts the regression line of $\log(\textit{incidence})$ on $\log \pi$ down, for fixed σ.

Breakout Theory, when $\sigma \approx 1$ and $\pi \approx 0$: Heterogeneity?

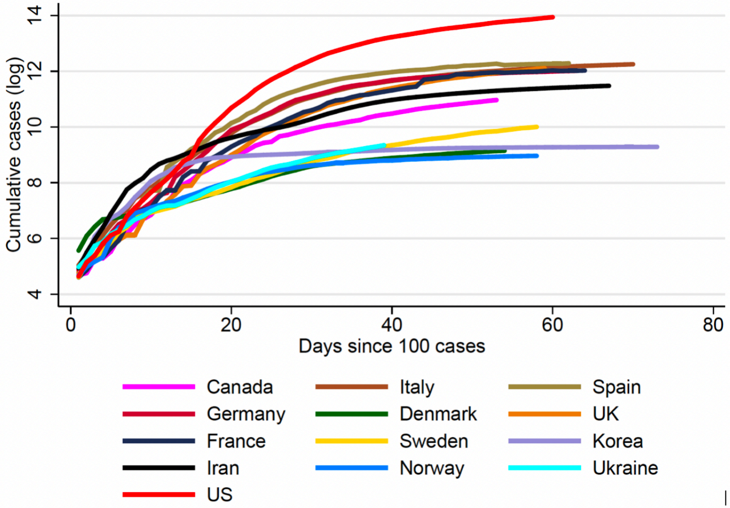

Sweden, Pre- and Post-Mitigation

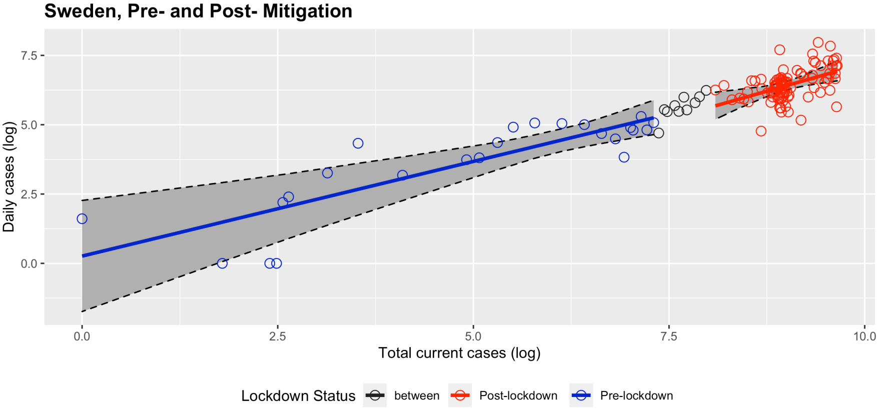

Pre-lockdown : log(Incidence) = 0.26 + 0.68 log(Prevalence), 95% CI slope: [0.25, 1.12]

Post-lockdown : log(Incidence) = 0.8 + 0.8 log(Prevalence), 95% CI slope: [0.48, 1.13]

UK, Pre- and Post-Lockdown

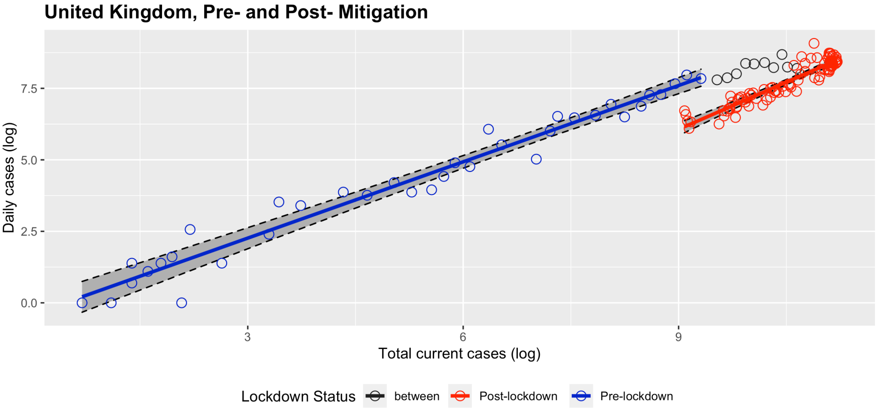

Pre-lockdown : log(Incidence) = -0.4 + 0.89 log(Prevalence), 95% CI slope: [0.83, 0.95]

Post-lockdown : log(Incidence) = 1.1 + 1.1 log(Prevalence), 95% CI slope: [1.01, 1.19]

France, Pre- and Post-Lockdown

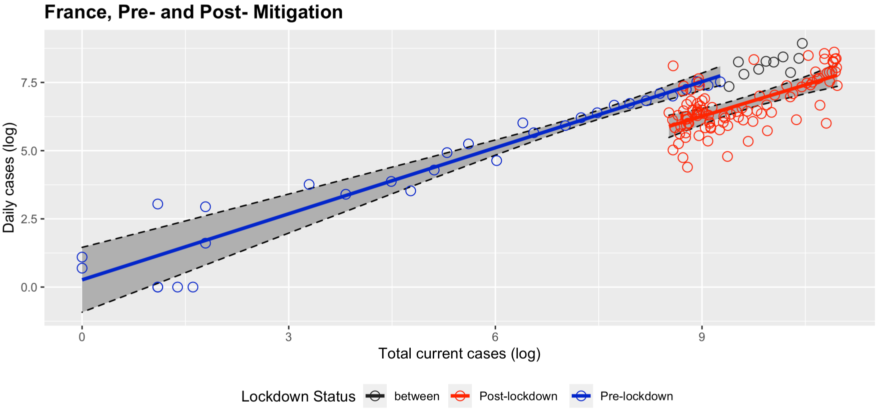

Pre-lockdown : log(Incidence) = 0.27 + 0.81 log(Prevalence), 95% CI slope: [0.68, 0.93]

Post-lockdown : log(Incidence) = 0.77 + 0.77 log(Prevalence), 95% CI slope: [0.6, 0.93]

Germany, Pre-Lockdown

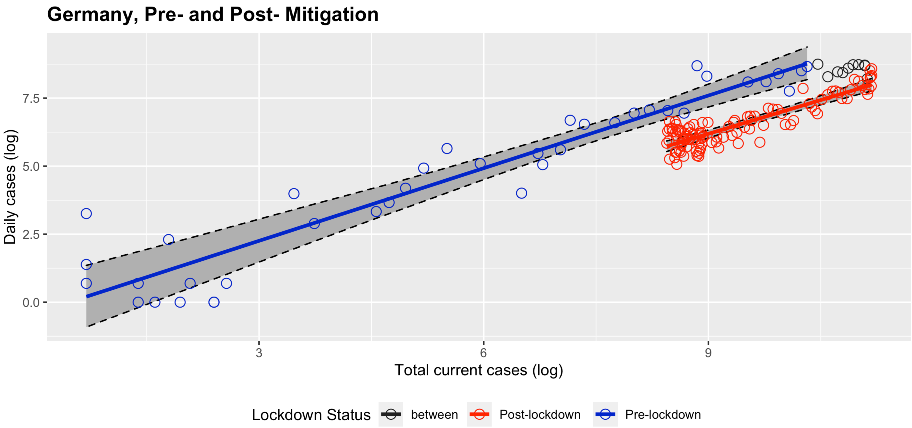

Pre-lockdown : log(Incidence) = -0.41 + 0.89 log(Prevalence), 95% CI slope: [0.77, 1.01]

Post-lockdown : log(Incidence) = 0.83 + 0.83 log(Prevalence), 95% CI slope: [0.75, 0.91]

Austria, Pre- and Post-Lockdown

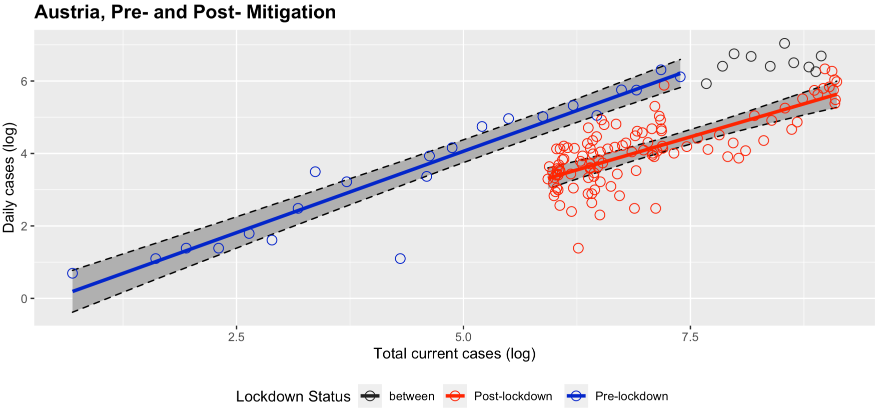

Pre-lockdown : log(Incidence) = -0.43 + 0.9 log(Prevalence), 95% CI slope: [0.81, 0.98]

Post-lockdown : log(Incidence) = 0.73 + 0.73 log(Prevalence), 95% CI slope: [0.63, 0.83]

Mexico, Pre- and Post-Mitigation

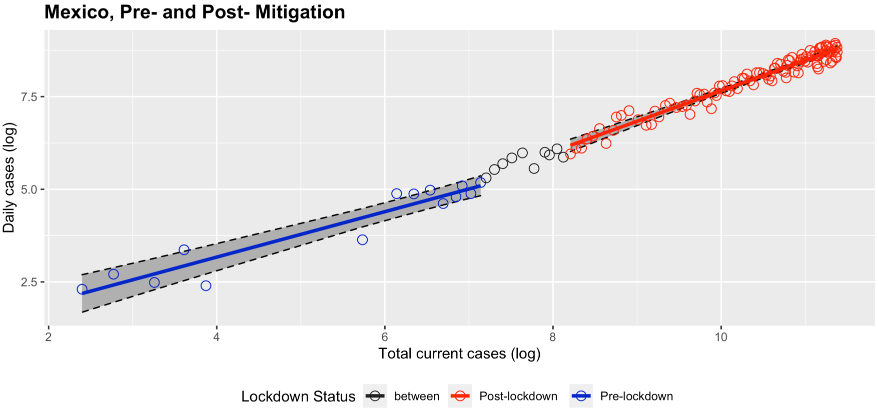

Pre-lockdown : log(Incidence) = 0.72 + 0.61 log(Prevalence), 95% CI slope: [0.51, 0.72]

Post-lockdown : log(Incidence) = 0.82 + 0.82 log(Prevalence), 95% CI slope: [0.78, 0.86]

Brazil, Pre- and Post-Mitigation

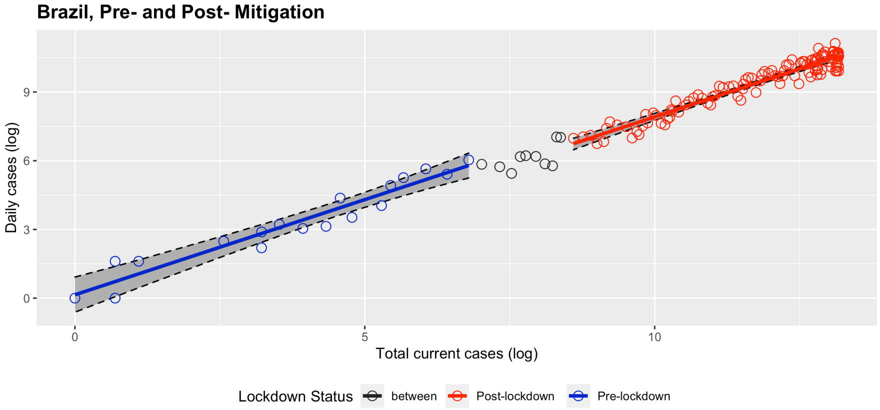

Pre-lockdown : log(Incidence) = 0.15 + 0.83 log(Prevalence), 95% CI slope: [0.68, 0.98]

Post-lockdown : log(Incidence) = 0.84 + 0.84 log(Prevalence), 95% CI slope: [0.79, 0.88]

Chile, Pre- and Post-Mitigation

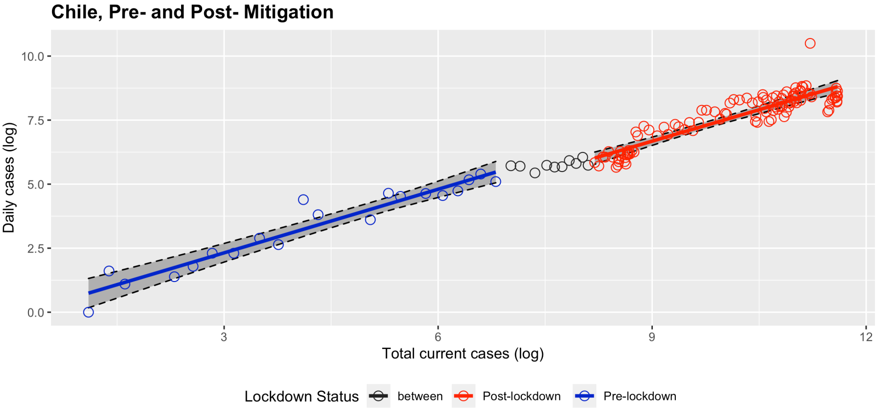

Pre-lockdown : log(Incidence) = -0.17 + 0.83 log(Prevalence), 95% CI slope: [0.69, 0.97]

Post-lockdown : log(Incidence) = 0.82 + 0.82 log(Prevalence), 95% CI slope: [0.75, 0.89]

New Zealand, Pre- and Post-Lockdown

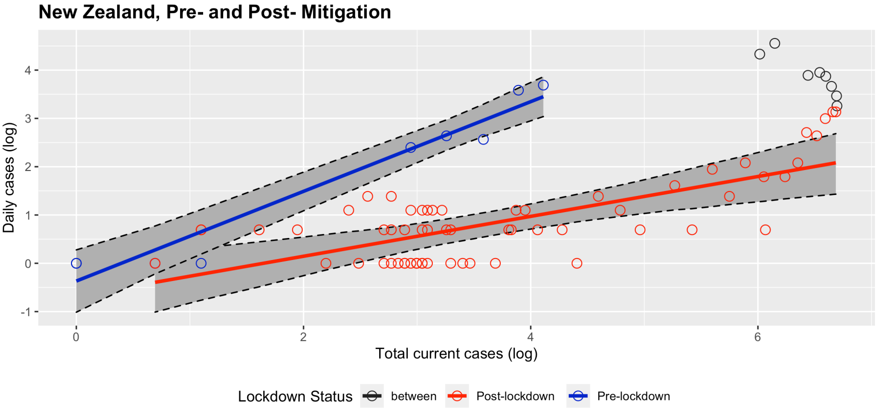

Pre-lockdown : log(Incidence) = -0.37 + 0.93 log(Prevalence), 95% CI slope: [0.64, 1.22]

Post-lockdown : log(Incidence) = 0.41 + 0.41 log(Prevalence), 95% CI slope: [0.28, 0.55]

New York, Pre- and Post-Lockdown

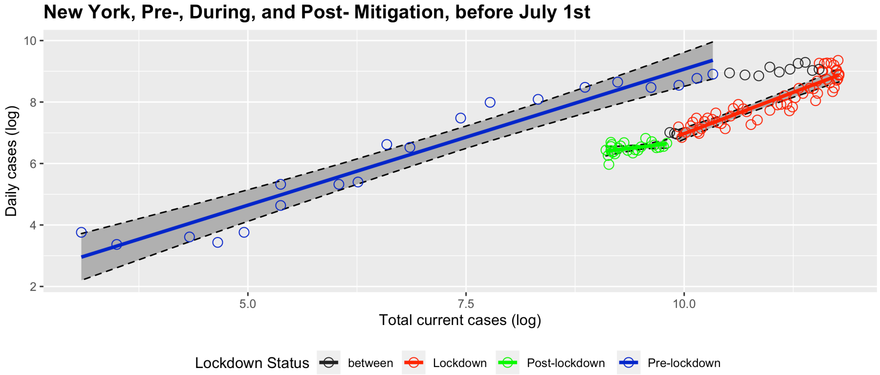

Pre-lockdown : log(Incidence) = 0.22 + 0.88 log(Prevalence), 95% CI slope: [0.75, 1.02]

Lockdown : log(Incidence) = 1.07 + 1.07 log(Prevalence), 95% CI slope: [0.96, 1.19]

Post-lockdown : log(Incidence) = 0.35 + 0.35 log(Prevalence), 95% CI slope: [0.08, 0.61]

Upwardly Biased Slope φ and Improving Testing

- Falling NY Positive-to-Test Ratio induces an omitted variable bias, that inflates the slope estimate φ

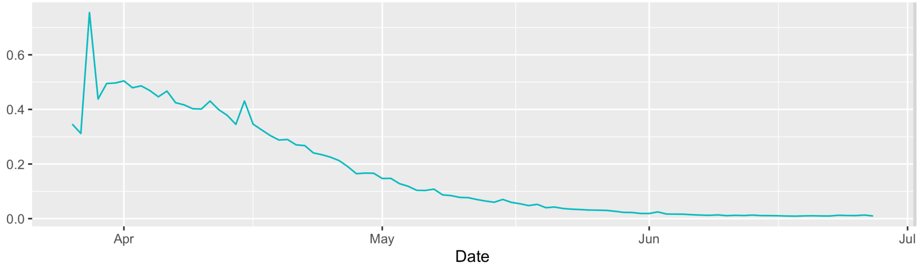

New Jersey, Pre- and Post-Lockdown

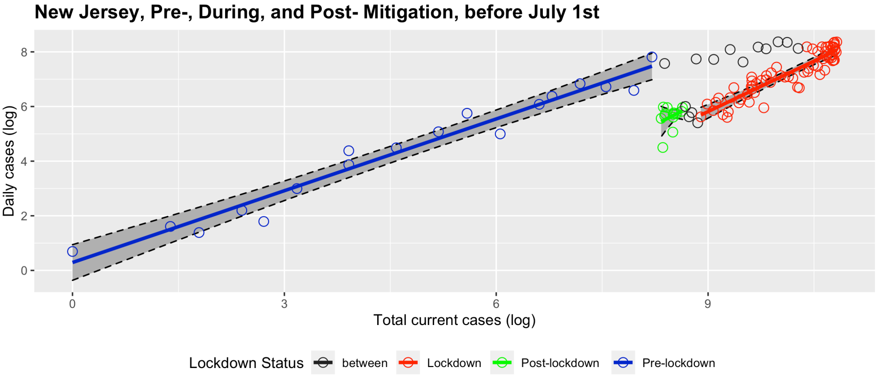

Pre-lockdown : log(Incidence) = 0.29 + 0.88 log(Prevalence), 95% CI slope: [0.76, 0.99]

Lockdown : log(Incidence) = 1.23 + 1.23 log(Prevalence), 95% CI slope: [1.12, 1.35]

Post-lockdown : log(Incidence) = 1.3 + 1.3 log(Prevalence), 95% CI slope: [-0.99, 3.59]

California, Pre- and Post-Lockdown/Mitigation

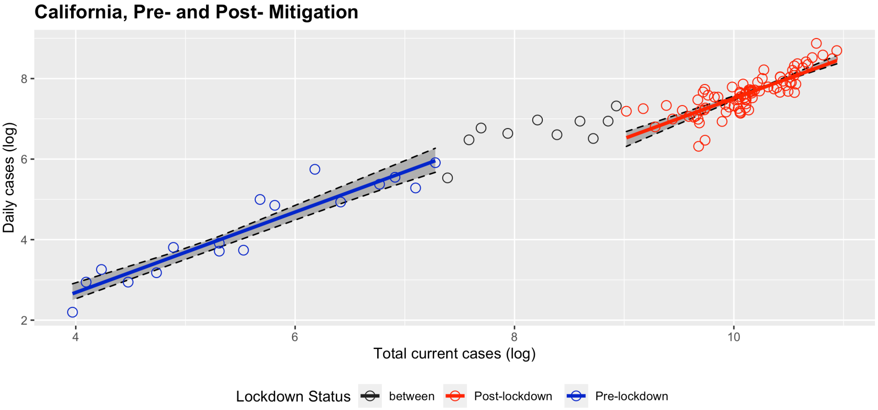

Pre-lockdown : log(Incidence) = -1.31 + 1 log(Prevalence), 95% CI slope: [0.81, 1.09]

Post-lockdown : log(Incidence) = 1 + 1 log(Prevalence), 95% CI slope: [0.85, 1.15]

Florida, Pre- and Post-Mitigation (Riots!)

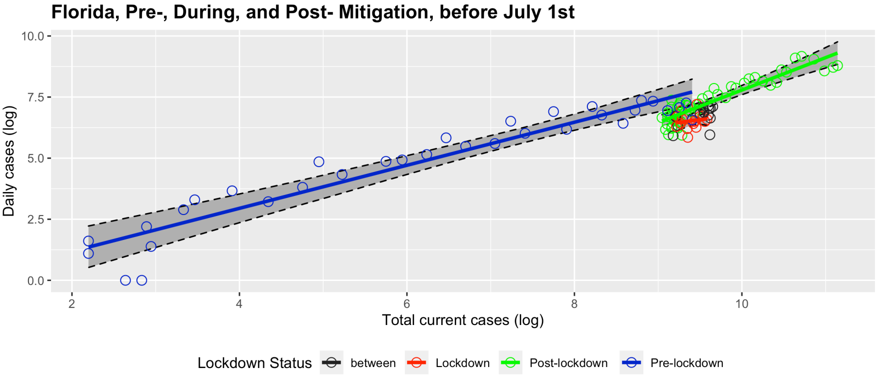

Pre-lockdown : log(Incidence) = -0.58 + 0.88 log(Prevalence), 95% CI slope: [0.75, 1.01]

Lockdown : log(Incidence) = 0.43 + 0.43 log(Prevalence), 95% CI slope: [-0.7, 1.56]

Post-lockdown : log(Incidence) = 1.33 + 1.33 log(Prevalence), 95% CI slope: [1.1, 1.56]

Texas, Pre- and Post-Mitigation (Riots!)

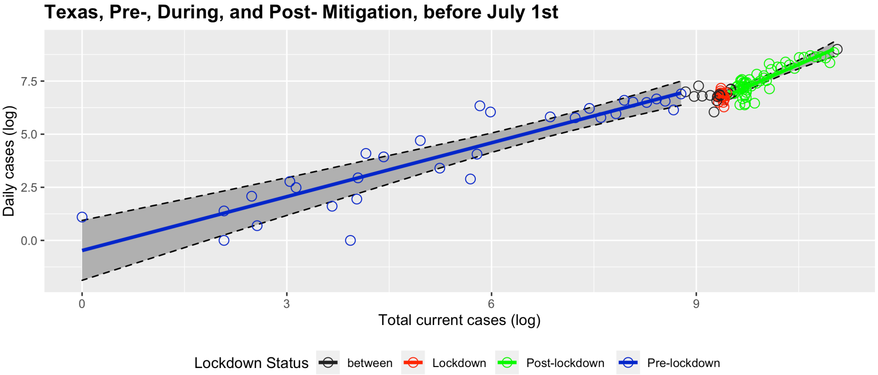

Pre-lockdown : log(Incidence) = -0.48 + 0.85 log(Prevalence), 95% CI slope: [0.65, 1.03]

Lockdown : log(Incidence) = -0.83 + -0.83 log(Prevalence), 95% CI slope: [-4.79, 3.15]

Post-lockdown : log(Incidence) = 1.4 + 1.4 log(Prevalence), 95% CI slope: [1.17, 1.62]

General Behavioral SIR Dynamics Nest the SIR Dynamics

- If $\pi_0 \leq \underline{\pi}$, SIR dynamics apply
- If $\pi_0 > \underline{\pi}$, then the vigilant regime starts. At this point:

$$\begin{aligned}\dot{\sigma}(t) &= -\beta q(\pi)\sigma(t)\underline{\pi}^{1-\varphi}\pi(t)^{\varphi} \\ \dot{\pi}(t) &= \beta q(\pi)\sigma(t)\underline{\pi}^{1-\varphi}\pi(t)^{\varphi} - r\pi(t)\end{aligned}$$

Theorem (Prevalence is Hump-Shaped)

In the BSIR, the susceptible share $\sigma(t)$ monotonically falls, while prevalence $\pi(t)$ either starts falling, or rises and then falls.

Famous Logarithmic Running Cases Plot

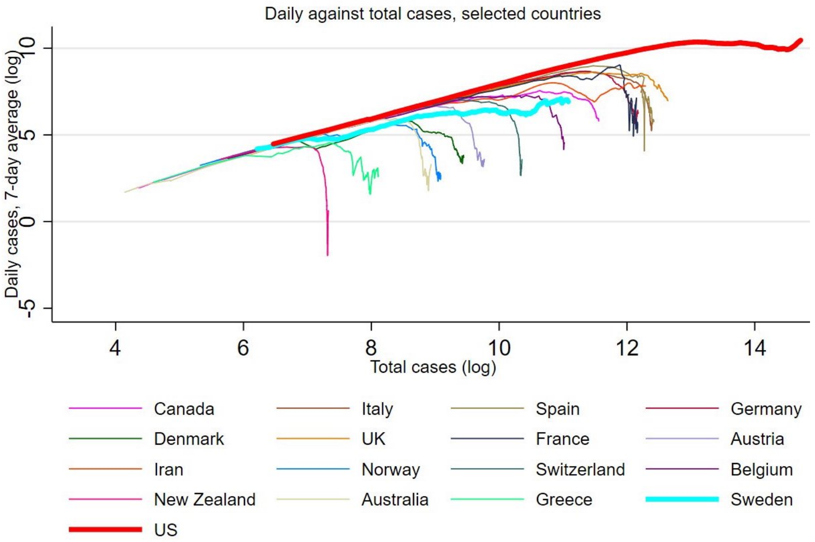

Breakout Theory

- SIR model: only immunity chokes off infections, and so bear breakout, log-linearity prevails
- For times $t < \tau$, the SIR dynamics apply:

$$\dot{\pi}(t) \approx \beta\pi(t) - r\pi(t) \quad \Rightarrow \quad \pi(t) \approx \pi_0 e^{(\beta-r)t}$$

- For times $t < \tau$, we have a Bernoulli differential equation:

$$\pi'(t) = \beta\underline{\pi}^{1-\varphi}\pi(t)^{\varphi} - r\pi(t) \quad \Rightarrow \quad \pi(t) = \underline{\pi}\left(\frac{\beta}{r}\left(1 - ke^{-r(1-\varphi)t}\right)\right)^{\frac{1}{1-\varphi}}$$

for the constant $k = (\beta/r - 1)\,(\underline{\pi}/\pi_0)^{r(1-\varphi)/(\beta-r)}$.

National Breakout Case Plots Over Time

- In the SIR model, these are log-linear.
- Assume a fraction α of non-spreading asymptomatics.

Theorem

Assume $\beta(1-\alpha) > r$.

- *In the chill regime (SIR model), π is increasing and log-linear.*
- *In the vigilant regime, prevalence $\pi(t)$ is increasing and logconcave, and is initially convex, eventually concave. Concavity happens sooner the lower is $\beta(1-\alpha)$ or ϕ.*

If $\beta(1-\alpha) < r$, then π is decreasing, logconcave, and convex.

Corollary

The sum of all past cases Υ is logconcave in time. It is convex when π is increasing, and concave when π is decreasing.

Mitigation or Lockdowns

- Think of mitigation or lockdown as a fall in the passing rate β.
- Here is a plot of $\pi(t)$ after β falls from 0.7 to 0.4.

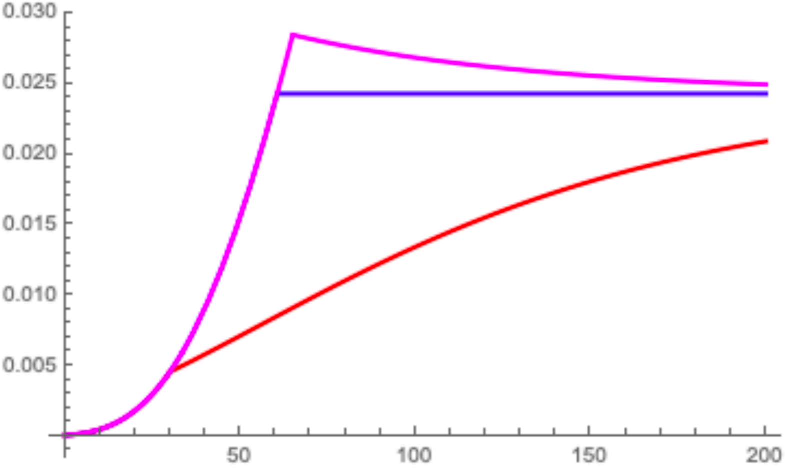

Herd Immunity

- Herd immunity tipping point:

$$B(t)\check{\sigma}_\varphi \check{\pi}_\varphi^\varphi = r\check{\pi}_\varphi \quad \Leftrightarrow \quad \check{\sigma}_\varphi = (r/B(t))\check{\pi}_\varphi^{1-\varphi} > r/\beta$$

Theorem

As the prevalence elasticity $\varphi \leq 1$ falls, (i) the herd immunity time τ_φ advances, (ii) the peak prevalence π_φ falls, (iii) the herd immunity infection share $1 - \sigma_\varphi$ falls, and (iv) its ratio to the eventual infection share $(1 - \sigma_\varphi)/(1 - \sigma_\varphi(\infty))$ rises.

The Road Ahead: SIR versus BSIR

- SIR Model: immunity chokes off contagions
- BSIR Model: immunity and vigilance choke off contagions

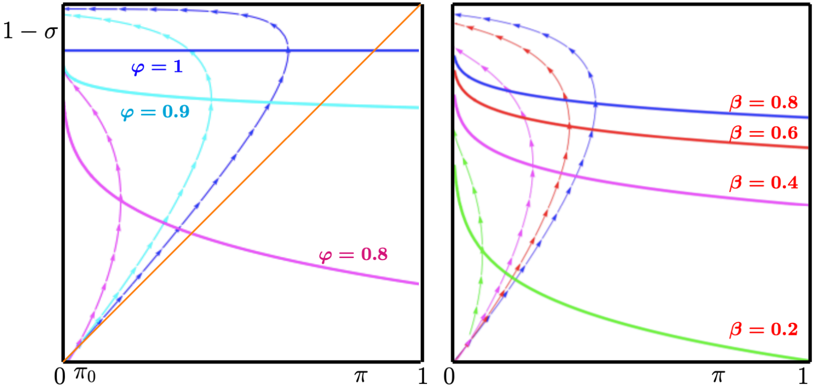

Herd Immunity — Behavioral SIR "Flattens the Curve"

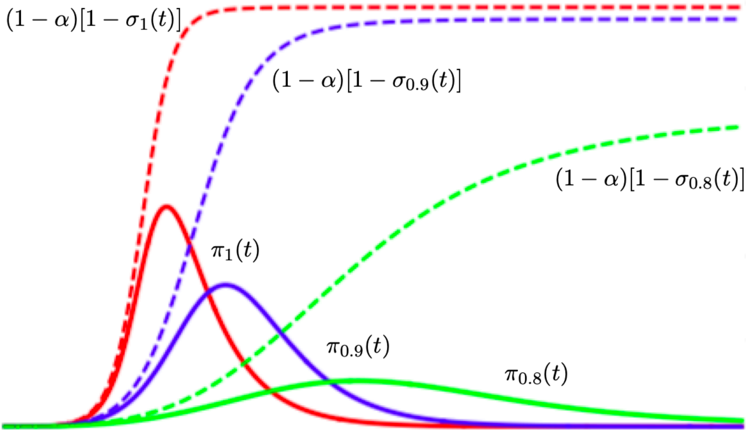

Herd Immunity Cases $\ll$ Eventual Total Cases

Eventual Infections

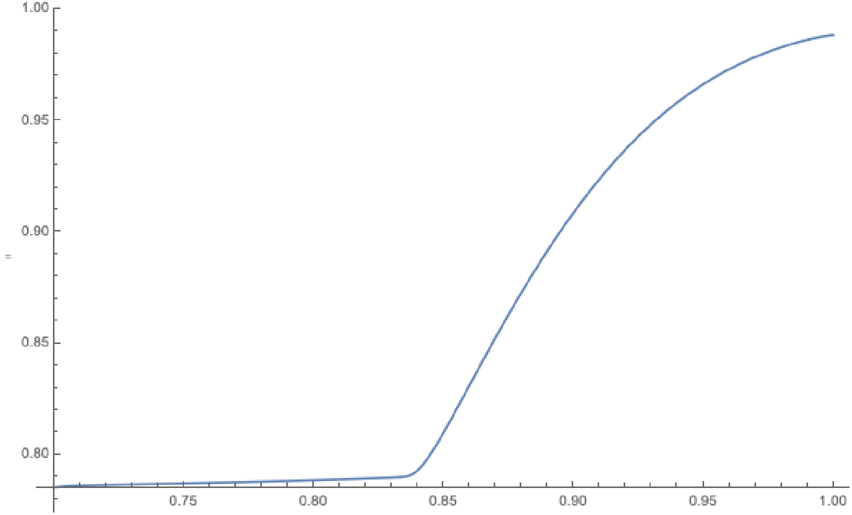

Herd Immunity Infections as a Share of Eventual Infections

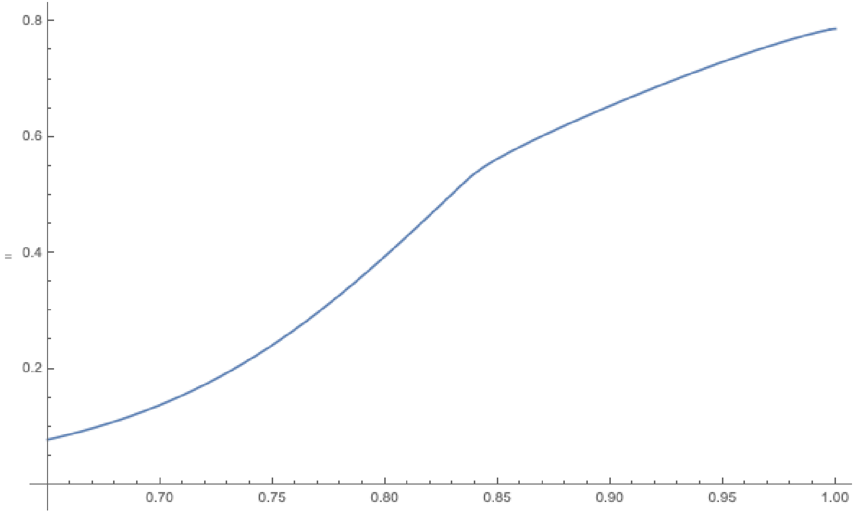

Swine Flu Herd Immunity

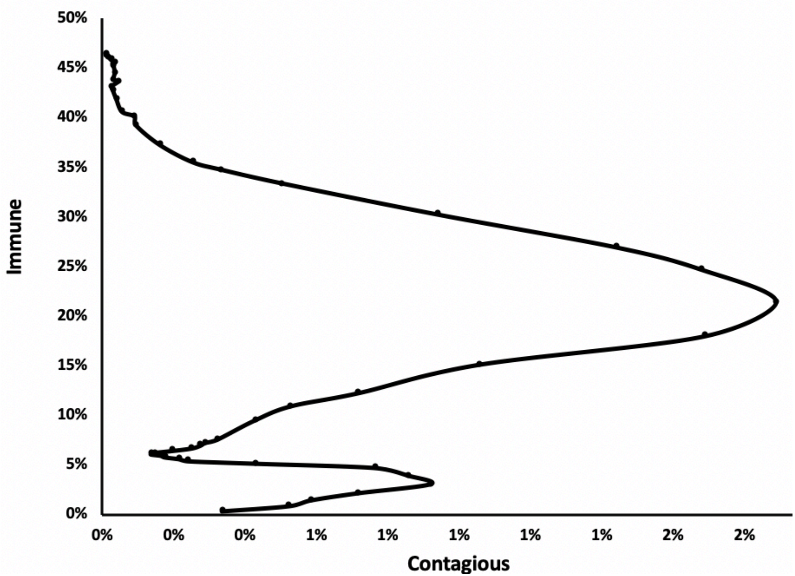

Swine Flu Herd Immunity

- Herd immunity on 10/31, 2009, with about 20% Immunity
- Lesson: about half of the sicknesses postdate herd immunity
- Lesson: the vaccine arrival in October was critical
- Lesson: seasonal component leads to “waves”

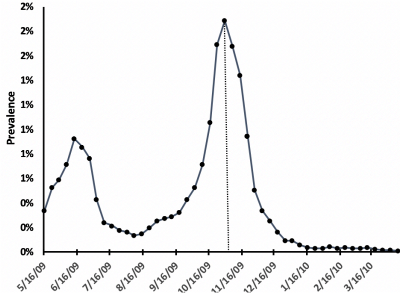

Swine Flu Herd Immunity