Optimal Epidemic Control in Equilibrium with Imperfect Reporting*

Thomas Phelan[1] Alexis Akira Toda[2]

[1]Federal Reserve Bank of Cleveland

[2]UCSD

March 27, 2021

*The views stated herein are those of the authors and are not necessarily those of the Federal Reserve Bank of Cleveland or the Board of Governors of the Federal Reserve System.

COVID-19 epidemic

- Started in December 2019 in Wuhan, China
- First community spread outside China and Korea reported in late February 2020 in Italy
- European countries quickly implemented mitigation policies such as lockdown
- U.S., Canada, and other countries followed suit by mid-March

Are mitigation policies warranted?

- At first glance, protecting public from infectious disease appears to be a responsible government action

Are mitigation policies warranted?

- At first glance, protecting public from infectious disease appears to be a responsible government action
- After second thought, draconian social distancing policies appears no longer self-evident
  1. Closure of schools and businesses have obvious economic costs
  2. Social distancing prevents new infections but also building *herd immunity*—hence policies may never achieve goal of ending epidemic ▸ VoxEU column
  3. Mathematical epidemic models used in policy-making ignore forward-looking behavior by rational agents
- Calls for rigorous theoretical (and quantitative) analysis

This paper

- Theoretically studies Susceptible-Infected-Recovered (SIR) epidemic model with forward-looking, rational agents
- Features:
  - Agents choose economic activity level optimally, understanding infection risk
  - Cases may be underreported and agents need to infer their health status (empirically, only 10–20% of cases reported)
  - Government may have limited commitment power to implement optimal policy

Main results (theoretical)

1. Prove existence of perfect Bayesian Markov competitive equilibrium
   - Here
     - "Perfect Bayesian": agents use Bayes rule on equilibrium path
     - "Markov": policy functions depend on state variables
     - "Competitive": agents ignore effects of their behavior on aggregate variables
   - Markov structure important because we can talk about optimal policy at any point in time or point in state space
2. Prove approximate static efficiency of equilibrium
   - Equilibrium inefficient due to externalities
     - Static: infected agents infect others
     - Dynamic: collective behavior of agents affect future dynamics
   - Static externality by unknown infected agents small

Main results (quantitative)

- Numerically solve for equilibrium and optimal policy of model of COVID-19 epidemic
- Findings:
  - Endogenous social distancing by agents mitigates epidemic, but welfare gain modest ($\sim 10\%$ reduction in welfare costs)
  - Quarantine effective, even with significant underreporting
  - Welfare gain from lockdown quite modest, and requires
    - Small enough number of initial cases
    - Underreporting
    - Rapid vaccine development
    - Government commitment until vaccine arrival

Literature

- **Mathematical epidemic model:** Kermack & McKendrick (1927)
- **Non-strategic economic epidemic model:** Sethi (1978), Kruse & Strack (2020)
- **Strategic epidemic model:**
  - (HIV) Geoffard & Philipson (1996), Kremer (1996), Auld (2003)
  - (Community infectious disease) Reluga (2010), Chen (2012), Fenichel (2013), Toxvaerd (2020)
- **Failure of well-intended public health policy** Toxvaerd (2019)

Society, agent types

- Agents: $n = 1, \ldots, N$ (finite but large)
- Time: $t = 0, \Delta, 2\Delta, \ldots$ (discrete, infinite horizon)
- Agent types by health/information
  - S: susceptible (no immunity)
  - I_k: known infected (reported)
  - I_u: unknown infected (unreported)
  - R_k: known recovered (reported)
  - R_u: unknown recovered (unreported)
  - D: dead
- Behavioral types are denoted by $h \in \{U, I_k, R_k, D\}$, where $U = S \cup I_u \cup R_u$ is unknown type

State variables

- Individual state variables are health/information status $h \in \{U, I_k, R_k, D\}$
- Aggregate state variables are fraction of each agent type
  - Let N_h be number of type $h \in \{S, I_k, I_u, R_k, R_u, D\}$ agents
  - With slight abuse of notation, let $h = N_h/N$ be fraction of type h agents
  - Then state space is set Z of $z = (S, I_k, I_u, R_k, R_u, D)$, where

$$S + I_k + I_u + R_k + R_u + D = 1,$$
$$N(S, I_k, I_u, R_k, R_u, D) \in \mathbb{Z}_+^6$$

- Assume aggregate state observable (can be inferred from random testing)

Preferences

- Alive agents (excluding I_k) take action $a \in A = [\underline{a}, 1] \subset [0, 1]$ with flow payoff $u(a)$, where $u(1) = 0$, $u' > 0$, $u'' < 0$
  - Interpretation: a is "economic activity level", with $a = 1$ being normal life and $a = \underline{a}$ being locked down
- Known infected (I_k) agents have utility function $u_I : A \to \mathbb{R}$, single-peaked at $a_I \in A$
- Dead agents receive flow utility $u_D < 0$, where $u_D \leq u_I(a) \leq u(a)$ for all $a \in A$
- Agents maximize expected discounted payoffs at rate $\mathrm{e}^{-r\Delta}$, where $r > 0$: discount rate

Disease transmission

- Agents meet randomly and transmit infectious disease
- If n, n' take actions a_n, a'_n, then they meet with probability $\lambda \Delta a_n a_{n'} / N$, where λ: meeting rate
  - λ exogenous and depends on how society is organized, e.g., population density, commuting, shopping, teaching pattern
- Conditional on n meeting n' and n' being infected, n gets infected with probability τ
  - τ exogenous and depends on contagiousness and society organization (greet by bowing, shaking hands, hugging, kissing)
- New infection reported with probability $\sigma \in (0, 1]$
- If type h agents take average action a_h, then infection probability

$$\Pr(\text{get infected} \mid a, \text{susceptible}) = \beta \Delta (a_{I_k} I_k + a_U I_u) a,$$

where $\beta = \tau \lambda$: transmission rate

Recovery, death, and vaccine arrival

- Every period, infected agents removed (recover or die) with probability $\gamma\Delta$
- I_u agents always recover; conditional on removal, I_k agents die with probability δ
  - δ: Case Fatality Rate (CFR) (fatality rate among reported cases)
  - $\delta_0 := \sigma\delta$: Infection Fatality Rate (IFR) (fatality rate among all cases)
- Recovered and dead agents remain so forever (lifelong immunity)
- Vaccine arrives at Poisson rate ν; once vaccine arrives, all alive non-infected agents become R_k

Mathematical epidemic models

- Our framework is basic SIR(D) model
- Other variations:
  - **SI:** infected agents remain infected forever (e.g., Epstein-Barr virus infection)
  - **SIS:** infected agents recover but can be reinfected (e.g., seasonal influenza)
  - **SEIR:** exposed agents are infected but not yet contagious
- For COVID-19, unclear whether immunity is lifelong
  - Probably SIRDS (possibility of $R \to S$ as in influenza) more appropriate, but in short run SIRD should be enough

Assumptions

Assumption (Perfect competition)

Agents view the evolution of the aggregate state z as exogenous and ignore the impact of their behavior on the aggregate state.

Assumption (Consistency)

On equilibrium paths, agents update their beliefs using the Bayes rule. Off equilibrium paths, unknown (U) agents believe they are susceptible with probability

$$\mu(z) := \begin{cases} \frac{S}{S+I_u+R_u} & \text{if } S > 0, \\ 0 & \text{otherwise.} \end{cases}$$

Individual problems: known recovered and dead

- Let $V_h(z)$ be value function of type h agents
- Because dead agents remain dead, Bellman equation is

$$V_D = (1 - \mathrm{e}^{-r\Delta})u_D + \mathrm{e}^{-r\Delta} V_D \iff V_D = u_D.$$

- Because known recovered agents remain recovered, Bellman equation is

$$V_{R_k} = \max_{a \in A} \left\{ (1 - \mathrm{e}^{-r\Delta})u(a) + \mathrm{e}^{-r\Delta} V_{R_k} \right\}.$$

- Optimal policy is clearly $a_{R_k} = 1$ and value function is $V_{R_k} = 0$

Individual problems: known infected

- Known infected agents are removed with probability $\gamma\Delta$, and conditional on removal, die with probability $\delta = \delta_0/\sigma$
- Hence Bellman equation is

$$V_{I_k} = \max_{a \in A}\{(1-\mathrm{e}^{-r\Delta})u_I(a) + \mathrm{e}^{-r\Delta}(\underbrace{(1-\gamma\Delta)V_{I_k}}_{\text{stay infected}} + \underbrace{\gamma\Delta[(1-\delta)V_{R_k} + \delta V_D]}_{\text{removal}})\}.$$

- Hence optimal policy is $a_{I_k} = a_I$ and value function is

$$V_{I_k} = \frac{(1-\mathrm{e}^{-r\Delta})u_I + \mathrm{e}^{-r\Delta}\gamma\Delta\delta u_D}{1-\mathrm{e}^{-r\Delta}(1-\gamma\Delta)},$$

where $u_I := u_I(a_I)$

Individual problems: unknown

- Suppose unknown agents adhere to policy function $a_U(z)$ and have belief $\mu(z)$ above
- Let $p(z) = \beta\Delta(a_I I_k + a_U(z) I_u)$ be probability of infection with full action ($a = 1$)
- Then Bellman equation is

$$V_U(z) = \max_{a \in A} \Bigg\{ (1 - \mathrm{e}^{-r\Delta}) u(a) + \underbrace{\mathrm{e}^{-r\Delta} \sigma \mu p a V_{I_k}}_{\text{known infection}}$$
$$+ \underbrace{\mathrm{e}^{-r\Delta}(1 - \sigma \mu p a)}_{\text{stay unknown}} \mathrm{E}_z(\underbrace{\mathrm{e}^{-\nu\Delta} V_U(z')}_{\text{no vaccine}} + \underbrace{(1 - \mathrm{e}^{-\nu\Delta}) V_{R_k}}_{\text{vaccine}}) \Bigg\}$$

Value functions

Proposition (Value functions)

Fix a policy function $a_U : Z \to A$ of unknown agents. Then there exists a unique value function $V_U : Z \to \mathbb{R}$ satisfying the Bellman equation. Furthermore,

$$\begin{aligned} V_D = u_D &< \frac{(\mathrm{e}^{r\Delta} - 1)u_I + \gamma\Delta\delta u_D}{\mathrm{e}^{r\Delta} - 1 + \gamma\Delta} = V_{I_k} \\ &< \frac{\mathrm{e}^{\nu\Delta}\sigma\beta\Delta}{\mathrm{e}^{(r+\nu)\Delta} - 1 + \sigma\beta\Delta} V_{I_k} \le V_U(z) \le V_{R_k} = 0. \end{aligned}$$

- $V_D \le V_{I_k}$ and $V_U \le V_{R_k}$ obvious
- $V_{I_k} \le V_U$ because U agents can always take $a = 1$ and $u_I \le u$

Equilibrium

- Our equilibrium concept is *perfect Bayesian Markov competitive equilibrium*
  - "Perfect Bayesian": agents use Bayes rule on equilibrium path
  - "Markov": policy functions depend on state variables
  - "Competitive": agents ignore effects of their behavior on aggregate variables
- Justification of competitive behavior: N large

Perfect Bayesian Markov competitive equilibrium

Definition (Markov equilibrium)

A (pure strategy) *perfect Bayesian Markov competitive equilibrium* consists of unknown agents' belief $\mu(z)$ of being susceptible, transition probabilities $\{q(z, z')\}_{z,z' \in Z}$ for aggregate state, value functions $\{V_h(z)\}_{h=U,I_k,R_k,D}$, and policy functions $\{a_h(z)\}_{h=U,I_k,R_k}$ such that

1. (Consistency) The belief $\mu(z)$ satisfies the Bayes rule on equilibrium paths; the transition probabilities $\{q(z, z')\}$ are consistent with individual actions and the mechanisms of disease transmission, symptom development, recovery, and death,
2. (Sequential rationality) Bellman equations hold

Existence of equilibrium

Theorem (Existence of equilibrium)

Under maintained assumptions, there exists a pure strategy perfect Bayesian Markov competitive equilibrium, where the belief $\mu(z)$ always satisfies

$$\mu(z) := \begin{cases} \frac{S}{S+I_u+R_u} & \text{if } S > 0, \\ 0 & \text{otherwise.} \end{cases}$$

Epidemic dynamics

- Fix any policy functions $(a_U(z), a_{I_k}(z))$ and induced transition probabilities $\{q(z, z')\}$,
- Then

$$\begin{aligned}
\mathrm{E}_z(S_{t+\Delta} - S_t)/\Delta &= -\beta a_U(z)(\sigma a_{I_k}(z) + (1-\sigma)a_U(z))S_t I_t, \\
\mathrm{E}_z(I_{t+\Delta} - I_t)/\Delta &= (\beta a_U(z)(\sigma a_{I_k}(z) + (1-\sigma)a_U(z)) - \gamma)I_t, \\
\mathrm{E}_z(R_{t+\Delta} - R_t)/\Delta &= \gamma(1-\delta)I_t, \\
\mathrm{E}_z(D_{t+\Delta} - D_t)/\Delta &= \gamma\delta I_t.
\end{aligned}$$

- Standard SIR(D) model special case by letting $a_U = a_{I_k} = 1$, $N \to \infty$, $\Delta \to 0$
- Definition: *herd immunity* achieved if $\mathcal{R}_0 S_t \leq 1$ (where $\mathcal{R}_0 = \beta/\gamma$: reproduction number), implying $\dot{I}_t \leq 0$

Model specification

- One period is a day
- Annual 5% discounting, so $r = 0.05/365.25$
- Vaccine arrives in one year, so $\nu = 1/365.25$
- Daily transmission rate $\beta = 1/5.4$, from meta-analysis of Rai et.al. (2021)
- Daily recovery rate $\gamma = 1/13.5$, from You et.al. (2020), so basic reproduction number $\mathcal{R}_0 = \beta/\gamma = 2.5$
- Infection fatality rate $\delta_0 = 0.0027$, from meta-analysis of Ioannidis (2021)
- Case fatality rate $\delta = 0.0135$, median value across 200+ countries/regions; hence reporting rate $\sigma = \delta_0/\delta = 0.2$

Model specification

- Log utility, so $u(a) = \log a$
- Assume $u_I(a) = u(a)$, so optimal action of I_k agents $a_I = 1$ (worst case analysis)
- Calibrate $u_D = -10.77$ from Sweden data, which did not introduce lockdown

Case fatality rate in Sweden

- $\delta_{\text{CFR}} = 0.0178$, so reporting rate $\sigma = 0.15$ in Sweden

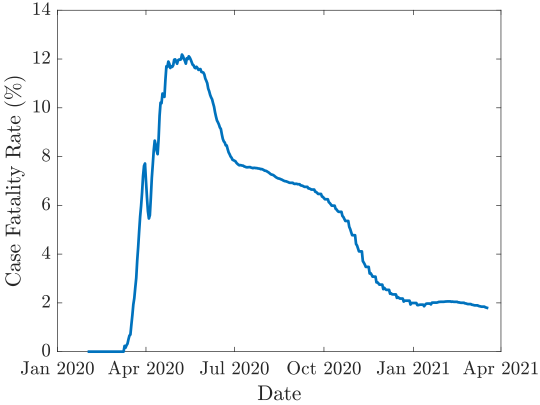

Prevalence in Sweden

- To back up prevalence, use accounting equation

$$D_{t+1} - D_t = \gamma\delta_{\mathrm{IFR}} I_t \iff I_t = (D_{t+1} - D_t)/\gamma\delta_{\mathrm{IFR}}$$

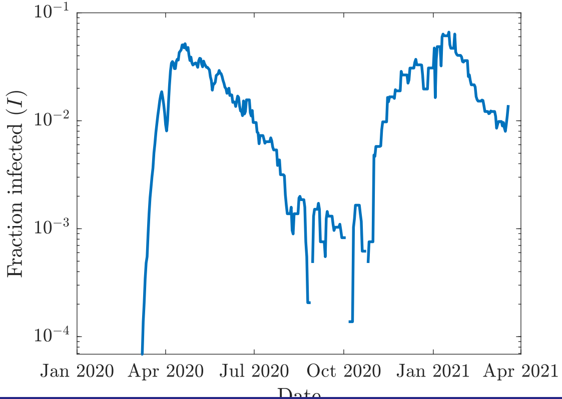

Myopic equilibrium

- Agents choose myopic optimal action $a = 1$ (standard SIR model)

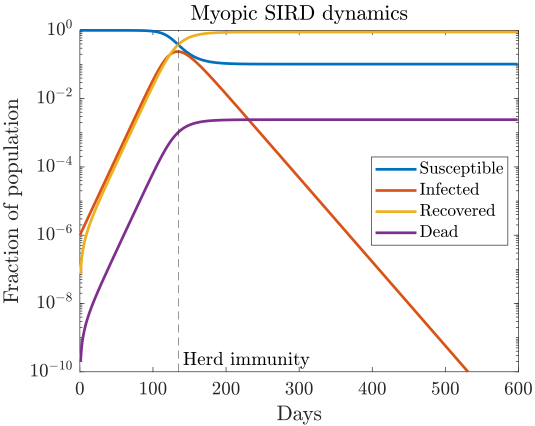

Markov equilibrium

- Agents choose individually optimal actions $a_U(z), a_{I_k}(z)$

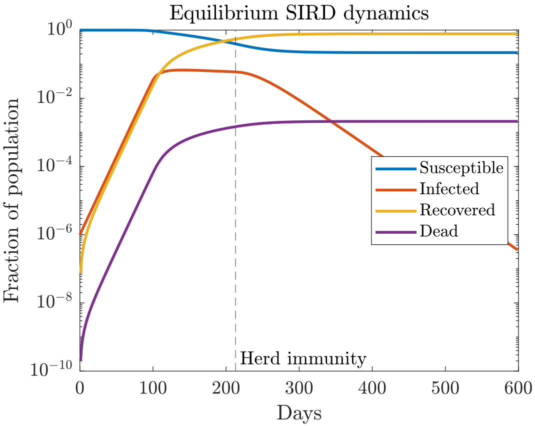

Equilibrium action of unknown agents

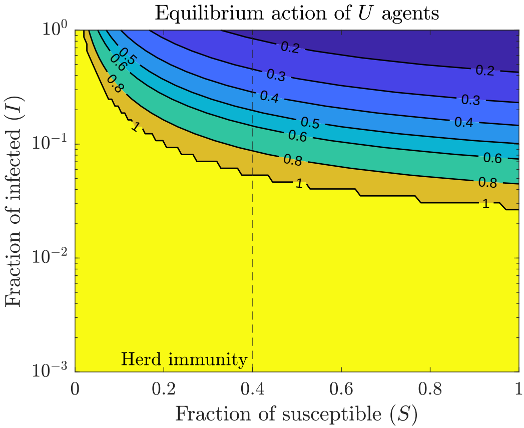

Efficient action of unknown agents

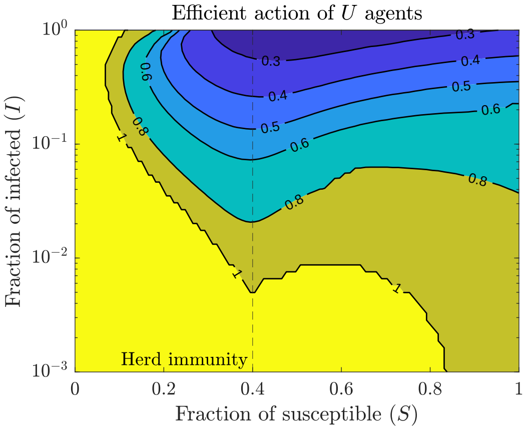

Recommended action over time

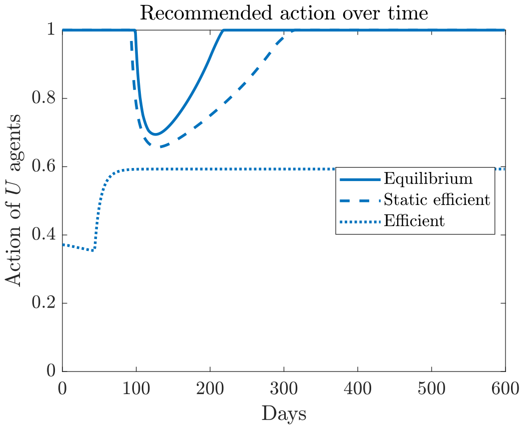

Welfare cost and death toll

- σ: reporting rate
- Welfare gains from lockdown quite modest

	Welfare cost (%)				Death toll (per 100,000)			
σ	My.	ME	SE	Eff.	My.	ME	SE	Eff.
0.1	1.77	1.61	1.58	1.34	242	212	200	0.0016
0.2	1.77	1.61	1.59	1.45	242	211	200	0.0016
0.4	1.77	1.62	1.61	1.56	242	209	200	183
0.7	1.77	1.62	1.61	1.57	242	207	202	183
1	1.77	1.60	1.60	1.55	242	203	203	182

Importance of quarantine

- So far we assume infected agents choose $a_I = 1$
- Suppose maximally quarantined, so $a_I = \underline{a}$

	Welfare cost (%)				Death toll (per 100,000)			
σ	My.	ME	SE	Eff.	My.	ME	SE	Eff.
0.1	1.77	1.42	1.36	1.11	242	203	189	0.0015
0.2	1.77	1.20	1.13	0.955	242	192	176	0.0015
0.4	1.77	0.616	0.554	0.584	242	158	139	73
0.7	1.77	0	0	0	242	0.0012	0.0012	0.0012
1	1.77	0	0	0	242	0	0	0

Efficient action without vaccine

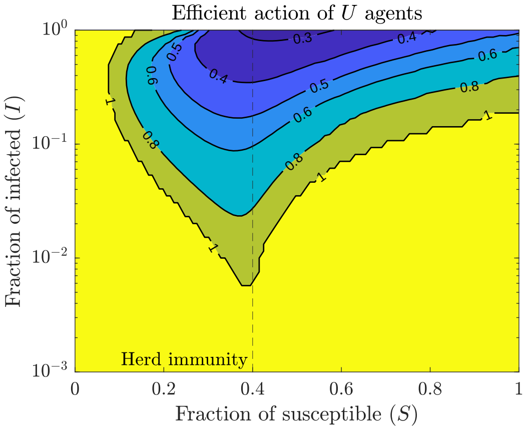

Recommended action without vaccine

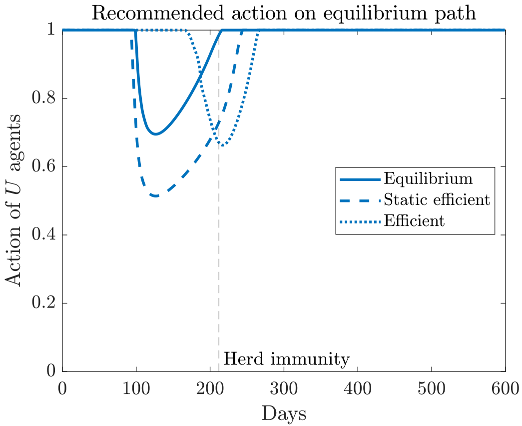

Welfare cost and death toll

σ	T_{vaccine}	Welfare cost (%)				Death toll (per 100,000)			
		My.	ME	SE	Eff.	My.	ME	SE	Eff.
0.2	1	1.77	1.61	1.59	1.45	242	211	200	0.0016
	1.5	1.99	1.86	1.89	1.70	242	211	200	185
	2	2.11	2.01	2.07	1.91	242	211	200	186
	5	2.34	2.30	2.43	2.15	242	211	200	189
	∞	2.51	2.51	2.72	2.32	242	211	200	190
1	1	1.77	1.60	1.60	1.55	242	203	203	182
	1.5	1.99	1.81	1.81	1.76	242	204	204	184
	2	2.11	1.93	1.93	1.87	242	204	204	185
	5	2.34	2.17	2.17	2.10	242	204	204	186
	∞	2.51	2.33	2.33	2.26	242	205	205	188

Importance of commitment power

- So far, we assume government have perfect commitment
- Assume now government fail to commit to optimal policy, and society reverts to equilibrium

	Welfare cost (%)			Death toll (per 100,000)		
T_{commit}	My.	ME	Eff.	My.	ME	Eff.
1/4	1.77	1.61	1.60	242	211	182
1/2	1.77	1.61	1.60	242	211	182
1	1.77	1.61	1.56	242	211	0.0016
2	1.77	1.61	1.52	242	211	0.0016
∞	1.77	1.61	1.45	242	211	0.0016

Conclusion

- Theoretically studied equilibrium model of epidemics with underreporting
- Findings:
  - Endogenous social distancing by agents mitigates epidemic, but welfare gain modest ($\sim 10\%$ reduction in welfare costs)
  - Quarantine effective, even with significant underreporting
  - Welfare gain from lockdown quite modest, and requires
    - Small enough number of initial cases
    - Underreporting
    - Rapid vaccine development
    - Government commitment until vaccine arrival