

Strategically Rational Risk Taking by Age in COVID-19, and the Heterogeneous Agent Behavioral SIR Model

Siamak Javadi* Elena Quercioli[†] Lones Smith[‡]

February 28, 2021

Abstract

Given the dramatic age variation in COVID death rates, we create a heterogeneous agent version of the Behavioral SIR contagion model of Engle et al. (2020). The Bayes Nash equilibrium of our infection avoidance game yields a simple new log-linear relationship between the case fatality rate (CFR) and COVID incidence: Everyone knows that everyone optimizes vigilance both for both the prevalence and their CFR.

We explain 2020 CDC incidence data for the USA north-east in terms of the CFR to age-specific COVID death data for Massachusetts. Our model is statistically significant: A 10% higher CFR reduces incidence by about 1%.

*Robert C. Vackar College of Business & Entrepreneurship, University of Texas at Rio Grande Valley. Email: siamak.javadi@utrgv.edu.

[†]Robert C. Vackar College of Business & Entrepreneurship, University of Texas at Rio Grande Valley. Email: elena.quercioli@utrgv.edu.

[‡]Economics Department, University of Wisconsin. Email: lones.smith@wisc.edu.

1 Introduction

A widely discussed feature story of the deadly COVID-19 pandemic has been the steep age-fatality profile: the youth rarely die, and the case fatality rate rises sharply with age. In fact,¹ adults over 85 have a death rate over 527 times that of adults age 18–29. This has led to immense differences in risk-taking behavior by age. This paper explains the falling infection rate in age as a strategically optimal “value-of-life” tradeoff between the risk of death from the pandemic and the costs of avoidance.

Engle et al. (2020) created a rational strategic twist on the classic SIR model² in which people minimize the sum of vigilance costs and expected disease losses. With a simple constant elasticity of avoidance in its costs, this game was fully solvable — namely, where everyone knows everyone else is optimizing too. The unique Nash equilibrium implied a log-linear map from prevalence to incidence with slope less than one. They statistically reject the nested SIR model for both the 2009 Swine Flu pandemic and COVID-19, finding an elasticity of incidence in prevalence significantly below one.

We enrich their *Behavioral SIR (BSIR) model* allowing for heterogeneous agents — like varying ages. In our Heterogeneous Agent BSIR Model, infection losses can vary across groups. It arises from the unique Bayesian Nash equilibrium of the multitype game. Incidence is not only log-linear in prevalence, with slope less than one — as in Engle et al. (2020) — but also log-linear in the infection loss, with negative slope. We find that optimization shaves about a tenth off marginal mortality changes. Specifically, we estimate our model using COVID-19 infection data from the CDC for the northeast USA, and case fatality rates (CFR) from deaths in Massachusetts. We find that a 10% increase in the CFR increases mortality by around 8.9%, after vigilance optimally adjusts. By contrast, the same 10% prevalence increase raises mortality only 8.4%.

By explicitly accounting for heterogeneous losses, we offer different evidence for the BSIR model. For the cross-sectional data allows us to identify three of its key strategic features: First, are youth behaving irrationally? Not in the economic sense: *Different age cohorts have maximized the identical objective functions, just with different losses.* Second, infection transmission reflects two parties’ avoidance efforts. Thus, a 10% prevalence rise increases one’s mortality risk that period less than a 10% increase in one’s CFR. For increased prevalence makes *both parties* to any meeting more vigilant, but a higher group CFR only impacts that group’s vigilance, and so typically impacts

¹See the CDC summary web page “COVID-19 Hospitalization and Death by Age”.

²SIR stands for “susceptible-infected-recovered” (Kermack and McKendrick, 1927)

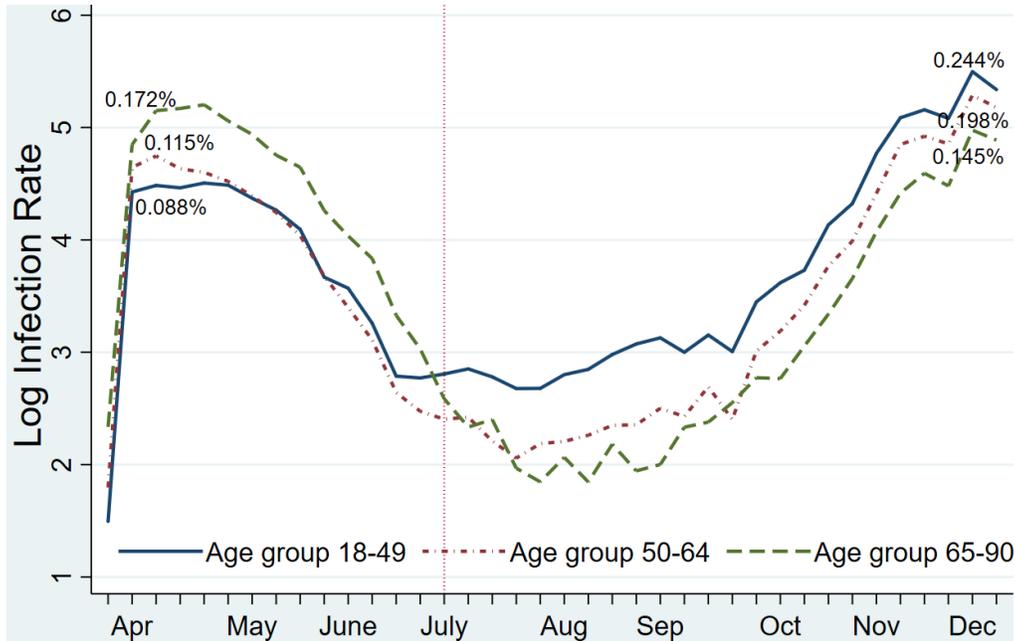


Figure 1: **The Infection Rate by Age Crossover (USA, North East).** COVID-19 infection rates in the northeast of the United States raged out of control in April and then fell. Focusing the magnifying glass on age groups, the rates were initially highest for the elderly, and lowest for the youth, and then switched.

one party’s vigilance. Third, this is a game of strategic substitutes: Higher vigilance by others depresses the marginal benefit of vigilance. As everyone is more vigilant when prevalence rises 10%, one does not increase vigilance as much as when one’s own CFR rises 10%. The last two properties can only be identified in a heterogenous agent model.

The average age of COVID infection fell over a decade just from May to August of 2020 (Boehmer et al., 2020). Figure 1 plots COVID-19 infections specifically for the USA North-East. The infection rates for our three age cohorts cross. While peak daily new COVID cases in July was more than twice that of April, the daily deaths after the July peak were only about two thirds of those in April (see worldometers.info/coronavirus) — reflecting the lower infection age.

McAdams (2020) summarizes the economic COVID literature, e.g. social distancing (Toxvaerd, 2020). Philipson and Posner (1995) first suggested the prevalence elasticity (of incidence); our context is a population game that yields a log-linear modification of the SIR model, and also yields an infection loss elasticity. Brotherhood et al. (2020) explore a macroeconomic model, with old and young individuals. Finally, with a dire future (not true here), forward-looking behavior can lead to fatalism (Auld, 2003).

Our theory in §2, 3, and A is self-contained. The empirical analysis is in §4.

2 The Model

In a large population, modeled as a unit continuum $[0, 1]$, everyone makes choices in a sequence of time periods (days or weeks). Fix a period. A mass $\sigma \in (0, 1)$ is initially *susceptible* to a disease, and $\pi \in (0, 1)$ is contagious; this is also called the *prevalence*.

The *SIR model* is the dynamical system for (π, σ) that assumes a flow *incidence* $I = \beta\sigma\pi$ of new infections, and some fixed exit rate $r > 0$ from the susceptible pool, either recovering or dying.³ This makes sense assuming random matching of identical individuals, and a fixed *passing rate* $\beta > 0$ when susceptible and infected persons meet.

The *Behavioral SIR model* (or BSIR) of Engle et al. (2020) modifies the SIR model, adding a costly *vigilance* action that reduces the passing rate below β .⁴ Examples of continuous vigilance include the fraction of time one wears a mask, or share of in-person meetings one skips. Vigilance $v \geq 0$ costs precisely v . If contagious and susceptible persons meet, with respective vigilances $v, w \geq 0$, *the passing rate is reduced to $\beta f(v)f(w)$* , for the *filter function* $f(v) = (v + 1)^{-\gamma} \in (0, 1]$, where $\gamma > 0$. In other words, vigilance reduces the passing rate, but with diminishing returns; infinite vigilance entirely chokes off infection: $f(0) = 1 > 0 = f(\infty)$. Also, the elasticity of f in “total vigilance” $v + 1$ is constant, namely, $-\gamma = (v + 1)f'(v)/f(v)$. So 1% vigilance increments result in constant $\gamma\%$ drops in the passing rate. It drives the log-linear formula in Theorem 1.

Anyone not sick or recovered thinks himself susceptible with probability $q \approx 1$.⁵ The *infection loss* $\ell > 0$ here varies by age group, unlike in the representative agent model of Engle et al. (2020). Let L denote the *random loss* for people one meets. As no one impacts future play, everyone must myopically best reply; therefore, equilibrium obtains for each period: Everyone minimizes the sum of vigilance costs and expected infection losses, solving

$$\min_v [v + \beta f(v)E[f(W)]q\pi\ell] \tag{1}$$

for any random vigilance W of others. In a *Bayes Nash equilibrium*, everyone solves optimization (1), and the random loss L yields random vigilance W as a best reply.

We assume that types mix uniformly and randomly — for random “out and about” encounters. The opposite extreme is that people sort by age, since it obviously formally reduces to the representative agent BSIR model, since all losses coincide.

³In continuous time, $\dot{\pi} = \beta\sigma\pi - r\pi$ and $\dot{\sigma} = -\beta\sigma\pi$. But we ignore these dynamics *in this paper*.

⁴As such, it applies to any infection model with a random transition from susceptible to infected, such as SIS (susceptible-infected-susceptible) and SI (susceptible-infected).

⁵If a small fraction $\omega \in (0, 1)$ of infected people is oblivious about it, then $q(\pi) = \sigma/(\sigma + \omega\pi) \approx 1$.

3 Equilibrium Analysis

Let $\mathcal{V}(\ell)$ be the *vigilance function*, namely, the best reply for someone with loss ℓ to others' vigilance $W \equiv V(L)$. Inspired by the derivation of the standard first price auction bidding function, when $V(\ell) > 0$, it solves a fixed point equation — namely, the equilibrium FOC, given random equilibrium vigilance $W \equiv V(L)$:

$$1 + f'(V(\ell))E[f(V(L))]\pi\ell\beta q = 0 \quad (2)$$

Barring a corner solution $V(\ell) = 0$ (ruled out in (★)), a solution of (2) is an optimum because $f'' > 0$ holds everywhere, guaranteeing the SOC and a unique interior solution. Since every type optimizes, such a vigilance function $V(\cdot)$ is a Bayes Nash equilibrium.

Lemma 1 (Equilibrium Filtering) *Define $C = \gamma\pi\beta q E[L^{-\gamma/(\gamma+1)}]$. In equilibrium, $f(V(\ell)) = C^{-\frac{\gamma}{2\gamma+1}}\ell^{-\gamma/(\gamma+1)}$ if $\ell \geq \ell_0$, while $f(V(\ell)) = 1$ if $\ell \leq \ell_0$, where $\ell_0 = C^{-\frac{1+\gamma}{2\gamma+1}}$.*

So above a threshold, greater losses ℓ lead people to filter out more infections. Our proof in the Appendix deduces the intuitive property that greater losses elicit more vigilance, so that the best reply vigilance $V(\ell)$ is increasing. Hence, the vigilance game has the *strategic substitutes* property: If others' vigilance $W = V(L)$ increases (in the sense of first order stochastic dominance), then the best reply $V(\ell)$ is lower.⁶

We assume a contagion so bad that everyone chooses positive vigilance — i.e., *the support of losses L is in $[\ell_0, \infty)$, for some $\ell_0 > 0$* (★). In what we will refer to as the *Heterogeneous Agent BSIR Model*, our next result applies Lemma 1 to derive a simple log-linear formula relating equilibrium incidence $I(\ell, \pi) = f(V(\ell))^2\beta\sigma\pi$ to ℓ and π .⁷

Theorem 1 (Incidence) *In the unique Bayes Nash equilibrium, given (★), the log incidence rate is*

$$\log I(\ell, \pi) = B + \psi \log \ell + \varphi \log \pi \quad (3)$$

for a constant B that depends on σ , β , L , γ , and q , and $\psi = \frac{-\gamma}{1+\gamma} < 0 < \varphi = \frac{1}{2\gamma+1} < 1$.

⁶Proof: For the expectation $E[f(V(L))]$ in the FOC (2) falls, since $|f'|$ is a decreasing function. To compensate, $f(V(\ell))$ must increase, and thus vigilance $V(\ell)$ must fall, since f is also decreasing.

⁷Theorem 1 implies simple dynamics for the heterogeneous agent BSIR. Specifically, there exists a threshold $\underline{\pi} \geq 0$ and an exponent $\varphi \in (0, 1)$, such that when $\pi \geq \underline{\pi}$, dynamics for $((\sigma_\ell), \pi)$ are:

$$\begin{aligned} \dot{\sigma}_\ell(t) &= -\beta\ell^\psi/E[L^\psi]q(\pi)\sigma_\ell(t)\underline{\pi}^{1-\varphi}\pi(t)^\varphi \\ \dot{\pi}(t) &= \beta[L^\psi\sigma_L(t)]/E[L^\psi]\underline{\pi}^{1-\varphi}\pi(t)^\varphi - r\pi(t) \end{aligned}$$

omitting a $q(\pi) \approx 1$ factor. We do not need these dynamics for the empirical analysis.

Theorem 1 applies to a mythical elegant continuum agent model. In §4, we estimate a significant result

$$\widehat{\log I} = B - 0.123(\log \ell) + 0.846 \log \pi$$

Slopes in (3) are standard economic elasticities. As in the homogenous agent BSIR model of Engle et al. (2020), the *incidence - prevalence elasticity* is $\varphi < 1$, rejecting the SIR model's assertion that $\varphi = 1$. Notably, this emerges in a panel regression here, rather than the time series analysis in Engle et al. (2020). So unlike the SIR model, incidence is not directly proportional to prevalence in the BSIR model. For vigilance adjusts to prevalence, shaving a constant fraction off infection rate changes.

The heterogeneity in losses yields our novelty: ψ , the *incidence - infection loss elasticity*. Theorem 1 has three testable takeout messages related to ψ that respectively reflect optimization, the matching externality, and the strategic substitutes property:

1. $\psi < 0$, since vigilance increases in the potential losses, just as it does prevalence. This effect simply reflects optimizing behavior.
2. $-\psi < 1 - \varphi$, since ψ reflects the vigilance efforts of just one party, but φ reflects the vigilance of both in this pairwise matching infection transmission model.
3. $1 - \varphi < -2\psi$, by the strategic substitutes property: people react less to greater prevalence than twice one person reacts to the same percentage CFR increase

The last two items follow from the explicit expressions for ψ and φ in Theorem 1.

For the second takeout message, note that Theorem 1 assumes independent random matching of types, where everyone selfishly optimizes, ignoring the positive externality of his vigilance. Namely, a 10% increase in a person's loss makes him more vigilant in all meetings, but the same 10% prevalence rise makes both parties more vigilant. At the opposite extreme, assume people sort by age. Given our optimization (1), a given percentage increase in the loss ℓ is formally equivalent to the same percentage increase in prevalence π ; therefore, the optimal solution for sorting by ages must be $\psi = \varphi - 1$. We will see that this is not the case in §4, that *infections are best understood as spread in random meetings rather than age-segregated encounters*.

4 Data and Empirical Analysis

To apply our theory, we assume that people are rational, and minimize the sum of vigilance and the expected loss of life. Let $\lambda > 0$ be the value of life, and Δ the probability of death, conditional on a COVID diagnosis.⁸ Then our loss in Theorem 1 is $\ell = \lambda\Delta$, given expected utility. Considering optimization (1), people are equally harmed in meetings by a 10% increase in the prevalence π and a 10% jump in the death rate Δ . Then the value of life λ is absorbed into the constant C in Lemma 1, along with $\kappa\pi$, and therefore the formula in Theorem 1 still applies, treating Δ as ℓ .

A. Probability of Death. To measure the probability of death Δ , we use the *case fatality rate* (CFR), or the share of COVID positives that end in death. We have found this data for one state, Massachusetts,⁹ which has so far had the third highest deaths per capita in the United States. This data includes hospital and nursing home deaths, which is important given what has transpired from late March to early August 2020, in age cohorts 0–19, 20–29, 30–39, 40–49, 50–59, 60–69, 70–79, 80+. To match with CDC data below, our paper will use the coarser partition of 0–19, 20–49, 50–69, 70+, the deaths were respectively 0, 146, 1210, and 7390, resulting in CFRs for these groups 0.00, 0.29, 3.71, 29.64 percent.¹⁰ We ignore the youngest age group, since it is not clear the extent to which they have rational independent agency, which our optimizing model requires. The death rates in this group are also too low to be accurately measured from our data (indeed, no deaths happen in our data).

We do not assume a constant CFR over time, as this is debated (Ledford, 2020). We instead create a piecewise linear time series of CFR’s, one for each age cohort i . We do so, using the time series of deaths in any week, divided by the new cases from three weeks earlier.¹¹ This reflects the variable time to die, centered about three weeks. In other words, we posit a constant rate of change b_i in each age group CFR_i , for $i = 1, 2, 3$, corresponding to youngest, middle aged, and elderly.

$$\log(CFR_{it}) = a_i + b_it + \epsilon_t$$

⁸The value of life is the willingness to pay for a small increment in the survival rate (Rosen, 1988).

⁹We use their COVID dashboard: www.mass.gov/info-details/covid-19-response-reporting

¹⁰Computations are based on positive COVID tests the week of March 22 through the week of July 19, and deaths through the week of Aug 9. Very few people died the first two weeks.

¹¹We ignore death underestimates, inferred from excess deaths, since this ratio does not vary greatly by age cohort, from 14.4% to 24.1% for all age cohorts over age 25 (Rossen et al., 2020).

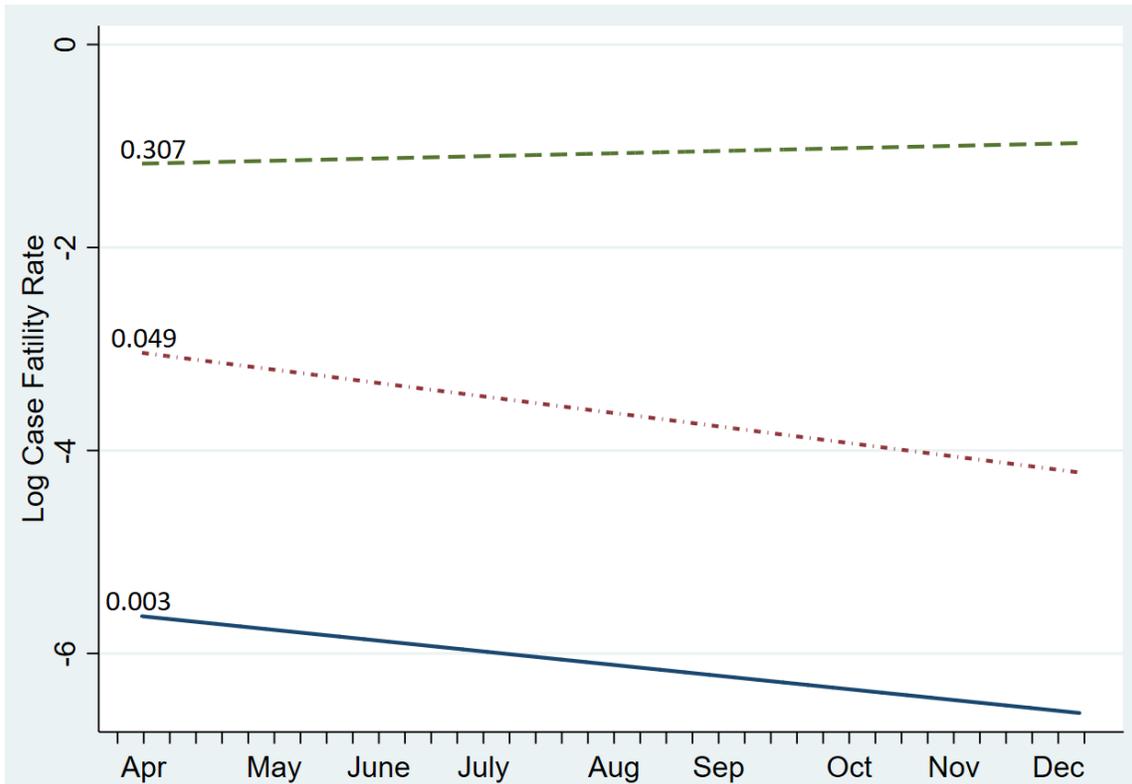


Figure 2: **The Estimated Log CFR with a Linear Time Trend.** We assume that individuals optimize on vigilance in response to their CFR. The age groups are 20-49 (solid) 50-69 (dot-dash), and 70+ (dash). The reported unlogged intercepts are roughly consistent with earlier noted CFRs. The slope estimates are not significant.

Using daily data for 21 weeks for each age group, we separately estimate the regressions depicted in Figure 2. This regression reflects how older individuals are more at risk from COVID, and medical treatments are evolving.

The CFR suffers from self-selection issues, as it conditions on a positive COVID test. This is known to be an undercount. The most reasonable interpretation of probability of death is the infection fatality rate (IFR). The IFR does not rely on testing, and thus can only be inferred from seroprevalence studies. But these studies only determine infections to date, and not the week by week infections. Using the CFR in lieu of the IFR is not unreasonable if the ratio of IFR to CFR does not vary greatly by age cohort. By weighting the year by year interpolated IFR's in Levin et al. (2020),¹² we compute the respective IFR's for these groups in Britain to be 0.060, 0.867, and 10.402.¹³

¹²Lacking such data for Massachusetts, we use age pyramid for France.

¹³The earlier noted raw CFRs 0.33, 4.12, 45.67 are resp. 5.5, 4.75, and 4.39 times bigger. This

CFR proxy \rightarrow	Entire sample	Excluding first 6 weeks
infection-loss elasticity γ	-0.123	-0.110
$P(H_0 : \gamma \geq 0)$	0.081	0.148
incidence-prevalence elasticity φ	0.846	0.837
$P(H_0 : -\gamma \geq 1 - \varphi)$	0.310	0.210
$P(H_0 : -\gamma \leq \frac{1}{2}(1 - \varphi))$	0.254	0.357
Number of Observations	111	96
Adjusted R^2	0.987	0.986

Table 1: **How Case Fatality Rates and Prevalence Impact Infection Rates.** The middle column documents the panel regression for all weeks CDC infection data, and the linearly projected CFR. The last column ignores the first six weeks when the nursing home deaths occurred.

B. Infection Rates. Finally, we turn to the COVID infection rates. To compute the infection rates by age, we use CDC data¹⁴ for HHS 1 region 1 (specifically, the states of CT, ME, MA, NH, RI, and VT). This contains the state of Massachusetts.

The CDC positive test rates were based on 6,419,892 specimens tested for SARS-CoV-2 using a molecular assay for the time span March 1–Dec 19, 2020. The percentage of specimens testing positive for SARS-CoV-2 each week, based on week of specimen collection, are summarized below. Unlike Massachusetts age groups, the CDC reports positive tests for 0-4, 5–17, 18–49, 50–64, and 65+. Merging the first two groups, the positive tests for 0-17, 18-49, 50-64, and 65+ are respectively, 6662, 50161, 32613, and 24931. We matched the CDC age groups to the amalgamated MA age groups, 20–49, 50–69, and 70+ respectively. This small mismatch is inescapable given the data coarseness limitations. Its impact is hopefully not major: It slightly increases the CFR of the middle age group, and slightly reduces the CFR of the oldest age group.¹⁵

In Table 1, we summarize two panel regressions with age and time fixed effects — one that ignores the first six weeks.¹⁶ This is the period when the vast majority of nursing home deaths occurred, before policy changes were enacted.

For COVID-19, people are maximally contagious from days 2–7. So inspired, it

consistency argues for the CFRs, since the testing does not appear to be too age-biased. *Notably, this also suggests that Massachusetts testing has undercounted COVID-19 by about a factor of five.*

¹⁴www.cdc.gov/coronavirus/2019-ncov/covid-data/covidview/01042021/specimens-tested.html.

¹⁵Using Levin et al. (2020), the interpolated infection fatality rates (IFR) per 100 for the MA cohorts is 0.0023, 0.06, 0.867, and 10.4, whereas the CDC cohorts has IFRs 0.019, 0.0567, 0.5895, and 7.93.

¹⁶We use a lag of three weeks, to match published estimates of time to due. To highlight the robustness of our findings, an online appendix shows that the estimated parameters are not too sensitive to how the CFR is computed, or whether we run this regression with daily data.

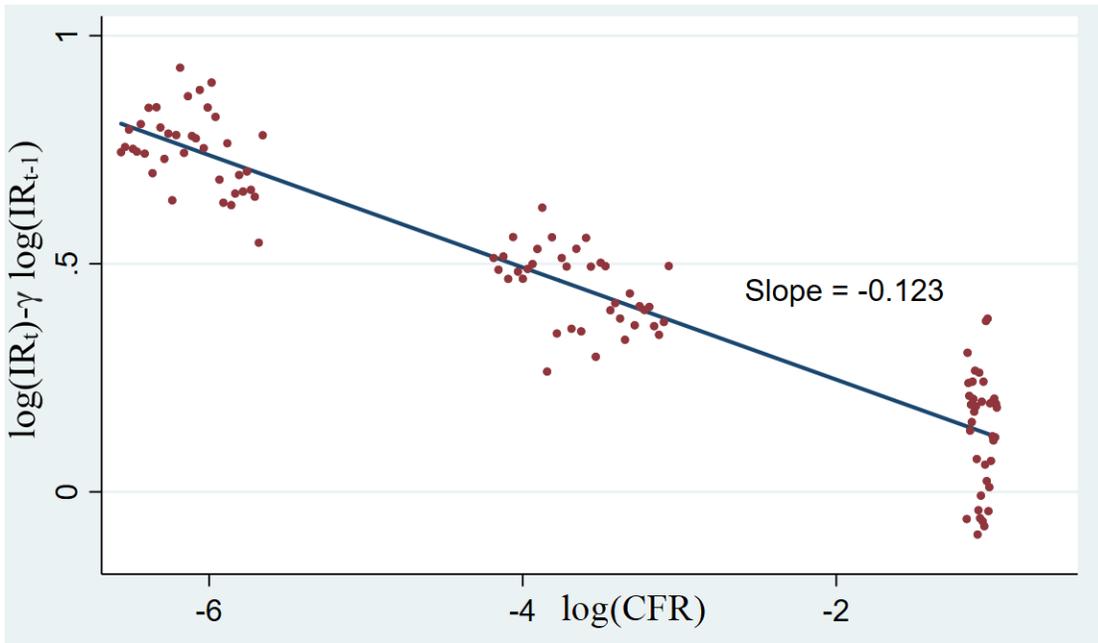


Figure 3: **The CFR Predicts the Infection Rate.** This is the scatter plot for the residuals of the regression of log infection rate on the CFR that are not explained by the last period infection rate, or by the age or time fixed effects.

is reasonable to proxy the prevalence by the percent infected last week IR_{t-1} , and incidence by the rate this week. Then Theorem 1 is proxied by the panel regression

$$IR_{i,t} = \alpha + \varphi IR_{i,t-1} + \gamma CFR_{i,t} + \delta_i + \tau_t + \varepsilon_{i,t}$$

for group effects δ_i and time effects τ_t . The time fixed effects capture the driving effect of prevalence on incidence that appears in the SIR and BSIR models; this panel approach avoids explicitly estimating the time series. This sidesteps any analysis of the time series aspects of the paper, which is the focus of Engle et al. (2020). Age effects capture heterogeneity in the age groups not summarized in the CFR. For instance, the youth may party more than other groups. Any undercounting of true infections, as long as it is same percent across ages, simply appears in the vertical intercept α .

Finally, let's consider our three predicted corollaries of Theorem 1 in §3. Consider first the optimizing behavior by agents solving the objective function (1). This involves two predictions, one old and one new. We estimate that incidence increases in prevalence with an elasticity $\hat{\varphi} < 1$, and this gap is significant. Specifically, a 10% increase in the prevalence leads to a 8.4% rise in incidence (not the 10% predicted by the SIR

model). This is consistent with the estimate found by Engle et al. (2020). And for our novel optimization as the case fatality rate changes, we estimate $\hat{\gamma} < 0$. In particular, we conclude that a 10% increase in a group’s CFR results depresses its incidence by 1.2%. In Figure 3, we give the scatter plots of $IR_t - \varphi IR_{t-1}$ to illustrate the impact of CFR on infection rates not explained by the prevalence, but captured by the CFR.

Consider finally the pairwise matching prediction $-\hat{\psi} < 1 - \hat{\varphi}$, namely, that two parties jointly respond more strongly to increased prevalence than does one to greater CFR, and the strategic substitutes prediction $1 - \hat{\varphi} < -2\hat{\psi}$, that efforts displace each other. These more nuanced predictions hold for our estimated parameters, but we cannot statistically reject the opposite inequalities. Both predictions await better data.

5 Conclusion

Three events must happen for death from COVID: exposure to the virus (prevalence), infection from the virus (passing), and death conditional on infection. Engle et al. (2020) introduce a strategic twist on the SIR model with endogenous vigilance. Their Nash equilibrium yields a simple log-linear map from prevalence to incidence that nests the SIR model as a special case. This paper reworks their avoidance game for a non-representative agent model, with a varying case fatality rate (CFR). Our unique (now) Bayesian Nash equilibrium yields a log-linear map from the CFR to incidence.

We show that this new model is predictive of cross-sectional behavior of different individuals in the pandemic. We deduce that avoidance behavior by each party shaves about 10% off the increased CFR by age: a 10% more deadly infection reduces the incidence by about 1%. Our parametric estimates are consistent with two other predictions of the model — the pairwise matching transmission or its strategic substitutes property. More refined data is needed to secure statistically significant tests of these.

Our paper and Engle et al. (2020) derive and test novel models of risk compensation with deadly consequences that apply matching theory and game theory. We could also compute willingness to pay, and so offer value of life analysis. Hopefully the COVID-19 pandemic soon ends, but our heterogeneous agent twist applies to any epidemic where people have divergent mortality risks, and equally well, divergent infection risks.¹⁷

¹⁷The latter holds for instance for HIV infections, with male to female transmission almost twice as likely as female to male transmission (ESGFHT-HIV, 1992).

A Appendix: Omitted Proofs

A.1 The Equilibrium Filtering Formula: Proof of Lemma 1

We now solve the differential equation (2) by further differentiating it. Since $f'(V(\ell))\ell$ is constant in ℓ by (2), its ℓ derivative is zero — and so $V(\ell)$ is differentiable. Hence:

$$f'(V(\ell)) + f''((V(\ell))V'(\ell)\ell) = 0 \quad (4)$$

As $f(v) = (v + 1)^{-\gamma}$ implies $f'(v)/f''(v) = -(v + 1)/(\gamma + 1)$, the equilibrium vigilance function solving (4) is

$$V(\ell) = c\ell^{1/(1+\gamma)} - 1 \quad (5)$$

All that remains is to solve for the constant $c > 0$. Vigilance vanishes at the loss $\ell_0 > 0$ where $V(\ell_0) = c\ell_0^{1/(\gamma+1)} - 1 = 0$. Then $\ell_0 = c^{-(1+\gamma)}$. In equilibrium, the filter function is then $f(V(\ell)) = (1 + V(\ell))^{-\gamma} = c^{-\gamma}\ell^{-\gamma/(\gamma+1)}$ if $\ell \geq \ell_0$, while $f(V(\ell)) = 1$ if $\ell \leq \ell_0$.

We now find the constant c . Since $\ell_0 = c^{-1-\gamma}$ by $V(\ell_0) = 0$ and (5), substituting $f'(V(\ell_0)) = f'(0) = -\gamma$, and the equilibrium filter function formula under (★), yield:

$$0 = 1 + f'(V(\ell_0))E[f(V(L))]\pi\ell_0\kappa = 1 - \gamma c^{-\gamma}E[L^{-\gamma/(\gamma+1)}]\pi\ell_0\kappa = 1 - \gamma\pi\kappa c^{-2\gamma-1}E[L^{-\gamma/(\gamma+1)}]$$

where $\kappa \equiv \beta q$. So the constant obeys $c^{2\gamma+1} = \gamma\pi\kappa E[L^{-\gamma/(\gamma+1)}]$. Now put $C = c^{2\gamma+1}$. \square

A.2 The Equilibrium Incidence Formula: Proof of Theorem 1

All told, if $\ell \geq \ell_0$, the equilibrium filter function is

$$f(V(\ell)) = c^{-\gamma}\ell^{-\gamma/(\gamma+1)} = \frac{\ell^{-\gamma/(\gamma+1)}}{(\gamma\pi\kappa E[L^{-\gamma/(\gamma+1)}])^{\gamma/(1+2\gamma)}}$$

Type ℓ 's incidence events (others to him, and him to others) $f(V(\ell))E[f(V(L))]\kappa\pi\sigma$ is:

$$\hat{I}(\ell, \pi, \sigma) = \ell^{-\gamma/(\gamma+1)}\pi^{1/(1+2\gamma)} \cdot E[L^{-\gamma/(\gamma+1)}]^{1/(1+2\gamma)}\gamma^{-2\gamma/(1+2\gamma)}\kappa^{1/(1+2\gamma)}\sigma \quad (6)$$

Then (6) yields (3) for interactions involving loss type ℓ . Finally, while these involve infection by ℓ of others, and vice versa, in a steady-state, the incidence density $I(\ell, \pi)$ of loss type ℓ is a constant fraction of the incidence events density $\hat{I}(\ell, \pi)$. Thus unknown fraction is absorbed in the constant B , along with the log of the last three factors (6).

References

- Auld, Christopher**, “Choices, Beliefs, and Infectious Disease Dynamics,” *Journal of Health Economics*, 2003, *22*, 361–377.
- Boehmer, T.K., J. DeVies, and et al. E. Caruso**, “Changing Age Distribution of the COVID-19 Pandemic — United States,” *MMWR and Morbidity and Mortality Weekly Report 2020*, 2020, *69* (39), 1404–09.
- Brotherhood, Luiz, Philipp Kircher, Cezar Santos, and Michèle Tertilt**, “An economic model of the Covid-19 epidemic: The importance of testing and age-specific policies,” 2020. CESifo Working Paper.
- Engle, Sam, Marianna Kudlyak, Jussi Keppo, Elena Quercioli, Lones Smith, and Andrea Wilson**, “The Behavioral SIR Model, with Applications to the Swine Flu and COVID-19 Pandemics,” 2020. mimeo.
- ESGFHT-HIV**, “Comparison of female to male and male to female transmission of HIV in 563 stable couples,” *BMJ*, 1992, *304* (6830), 809–13. European Study Group’s findings on Heterosexual Transmission of HIV.
- Kermack, William Ogilvy and A. G. McKendrick**, “A Contribution to the Mathematical Theory of Epidemics,” *Proceedings of the Royal Society*, 1927, *115*, 700–721.
- Ledford, Heidi**, “Why do COVID death rates seem to be falling?,” *Nature*, November 2020, pp. 190–192.
- Levin, Andrew T., William P. Hanage, Nana Owusu-Boaitey, Kensington B. Cochran, Seamus P. Walsh, and Gideon Meyerowitz-Katz**, “Assessing the age specificity of infection fatality rates for COVID-19: systematic review, meta-analysis, and public policy implications,” *European Journal of Epidemiology*, 2020, *35*, 1123–1138.
- McAdams, David**, “Economic epidemiology in the wake of Covid-19,” *Covid Economics*, 2020, pp. 1–45.
- Philipson, Tomas and Richard A Posner**, “The microeconomics of the AIDS epidemic in Africa,” *Population and Development Review*, 1995, pp. 835–848.

Rosen, Sherwin, “The value of changes in life expectancy,” *Journal of Risk and Uncertainty*, 1988, *1*, 285–304.

Rossen, LM, AM Branum, FB Ahmad, P Sutton, and RN Anderson, “Excess Deaths Associated with COVID-19, by Age and Race and Ethnicity,” *MMWR and Morbidity and Mortality Weekly Report 2020*, 2020, *69* (42), 1522–1527.

Toxvaerd, Flavio, “Equilibrium Social Distancing,” 2020. March mimeo, Cambridge.