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

Siamak Javadi* Elena Quercioli† Lones Smith‡

May 30, 2023

Abstract

Given the dramatic age variation in COVID death rates, we create a heterogeneous agent version of the Behavioral SI* contagion model of Keppo et al. (2020). Individuals randomly meet each others pairwise, unaware of their types. Inspired by auction theory, we compute the Bayes Nash equilibrium of the pairwise incomplete information games transpiring over time and across the population. This 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.

†*Elena sadly died in summer 2021*, while on tenure track in Robert C. Vackar College of Business & Entrepreneurship, at University of Texas Rio Grande Valley. :(Lones dedicates this paper to her, which reflected her our last insights on our behavioral SI* model.

‡Economics Department, University of Wisconsin. Email: lones.smith@wisc.edu. I am grateful to the National Science Foundation for funding this research.

1 Introduction

A widely discussed feature 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.

We build on Keppo et al. (2020), who produce a new and solvable game that captures a strategic pairwise version of the classic *Swiss Cheese Model* of accident causation in random encounter settings — where accidents only arise from multiple aligned failures. In a contagion setting, this corresponds to someone passing along a virus, and another failing to guard against it: it passes though both players' filters. In their optimizing twist on the classic SI* model[2], the infection passing rate reflects costs and benefits. The costs are an additive cost of *vigilance*, while the benefits come from the lower passing rate. People are fully rational, and minimize the sum of vigilance costs and expected disease losses. By assuming a constant elasticity of avoidance in vigilance costs, this game was fully solvable. In the unique Nash equilibrium, everyone knows everyone else is optimizing too, and understands their vigilance choice. It 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.

Our model starts with the observation that if vigilance is motivated by the fear of death, then doubling the death rate conditional on infection yields the same risky tradeoffs as doubling the prevalence rate. As a result, the same equilibrium substitution effect in the representative agent analysis in Keppo et al. (2020) should intuitively emerge when people vary by conditional death rate. We therefore enrich their *Behavioral SI* model* allowing for heterogeneous infection losses agents — such as from varying ages. We derive a unique Bayesian Nash equilibrium of the multitype game — which reduces to the unique Nash equilibrium with only one type. Incidence is now not only log-linear in prevalence with slope less than one, but also log-linear in the infection loss, with negative slope.

[1] See the CDC summary web page "COVID-19 Hospitalization and Death by Age".

[2] Here SI stands for "susceptible-infected". applies to any infection model with a random transition from susceptible to infected, such as SIS (susceptible-infected-susceptible) and SI (susceptible-infected). So SIR stands for "susceptible-infected-recovered" (Kermack and McKendrick, 1927).

We estimate our model using COVID-19 infection data from the CDC for the north-east 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%. In other words, optimization shaves about a tenth off marginal mortality changes.

By explicitly accounting for heterogeneous losses, we offer different evidence for the model in Keppo et al. (2020). For the cross-sectional data allows us to identify three of its key strategic features: First, did the youth behaving irrationally in the COVID pandemic? Not in the economic sense: *Different age cohorts have maximized the identical objective functions, just with different losses.* Their greater infection levels are rationally thus, reflecting their lower losses in the quasilinear optimization.

Second, infection transmission reflects avoidance efforts by two parties. Thus, a 10% higher prevalence increases one's mortality risk that period less than a 10% increase in one's CFR. For prevalence impacts *both parties'* vigilance in any meeting, but a higher CFR of some type only impacts that type's vigilance. This explains why the fantastic escalation in death rates with age is not reflected by the same percentage reduction in infection rates as a similar rise in COVID prevalence.

Third, our random encounter game exhibits strategic substitutes. For since one's filter also protects others if one is infected, greater vigilance by others depresses the marginal benefit of vigilance — reducing its optimal level. Since everyone is more vigilant when prevalence rises 10%, vigilance should not increase as much as when one's own CFR rises 10%. Notably, the last two properties can now be identified in this heterogenous agent model, and we in fact confirm them in the data.

The above findings crucially rely on individuals not sorting by age in their random encounters, If that happened, then the changes in CFR by age would be reflected in one's partner's actions changing likewise, and we would not see these latter two findings. This highlights a new theoretical reason why the pandemic was spread by mixing across ages.

Consistent with individuals re-optimizing in response to advertised death rates, 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

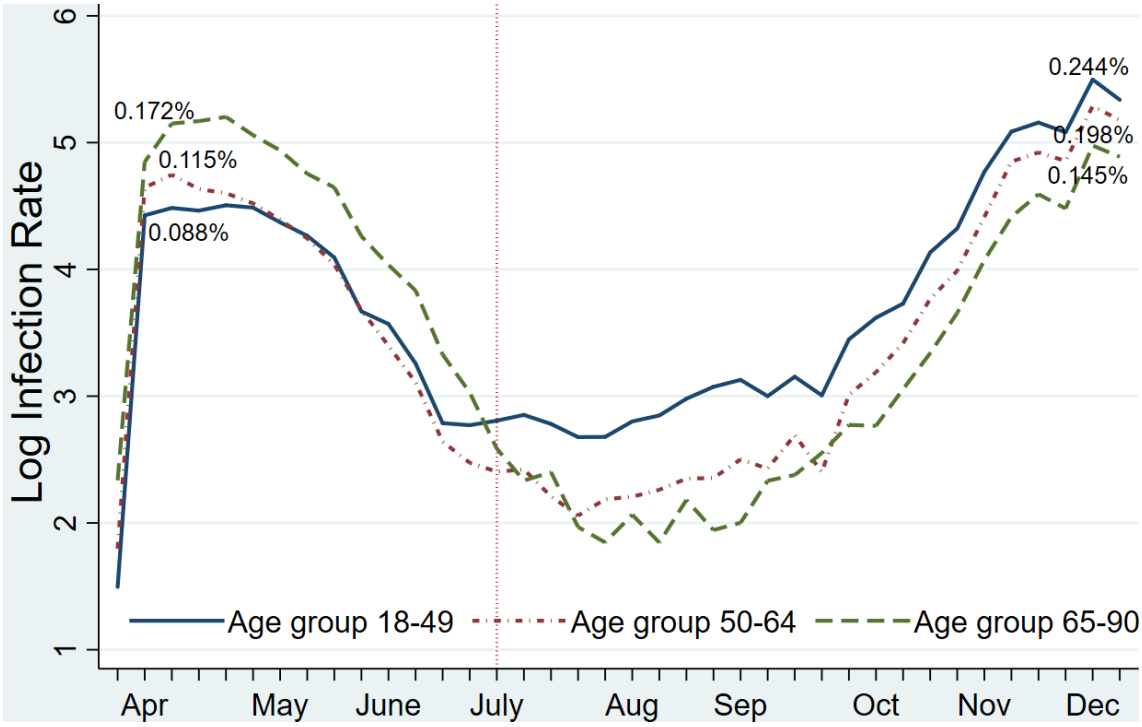

Figure 1: **The Infection Rate by Age Crossover (USA, North East).** COVID-19 infection rates in the northeast of the United States raged out of control in April and then fell. Focusing the magnifying glass on age groups, the rates were initially highest for the elderly, and lowest for the youth, and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

peak were only about two thirds of those in April[3]— reflecting the lower infection age.

McAdams (2020) summarizes the economic COVID literature, e.g. social distancing (Toxvaerd, 2020).[4] 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).

Compared to Keppo et al. (2020), our solution method is new. We take inspiration from the differential equation approach to revelation principle that is used to solve for the bidding strategy in the unique Bayesian Nash equilibrium of first price auction. Here, instead of asset values differ, the infection loss varies.

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

[3]See worldometers.info/coronavirus.

[4]Examples of continuous vigilance for epidemics include the fraction of time one wears a mask, or the share of face-to-face meetings one skips.

2 The Model

In a large population modeled as a unit mass continuum $[0, 1]$, everyone makes choices in a sequence of time periods (e.g. days). Each period, some mass $\sigma \in (0, 1)$ is *susceptible* to a disease, while the *prevalence* is the contagious mass $\pi \in (0, 1)$. Both π and σ are commonly known. In the SIR model, the *incidence* of new infections is $I = \beta\sigma\pi$ when susceptible and infected persons meet. Here, $\beta > 0$ is a fixed *passing rate*. This product structure follows from random encounters, independent of infection status.

In the behavioral twist on the SIR model by Keppo et al. (2020), infection passing reflects costs and benefits in each period. They assume that costly real-valued *vigilance* action — like time spent wearing a mask — can reduce the passing rate below its baseline β. Label the vigilance action by its cost $v \geq 0$. If a contagious person meets someone susceptible, while exerting vigilances $v, w \geq 0$, then the passing rate falls to $\beta f(v)f(w)$. Keppo et al. (2020) posit $f(v) = (v + 1)^{-\gamma} \in (0, 1]$, for which $f' < 0 < f''$ and $f(0) = 1 > 0 = f(\infty)$. This yields a constant elasticity of the passing rate in "total vigilance" $v + 1$. So 1% more total vigilance lowers the passing rate by γ%.

To capture how older people died from COVID-19 far more often, we modify the representative agent model of Keppo et al. (2020) so that people one meets vary in their *infection loss* $\ell > 0$. We show how behavior varies in equilibrium.

Let L denote the *random loss* for people one meets. Let $q \approx 1$ be the chance that any individual is susceptible who has not yet been symptomatically infected. Someone uninfected of loss ℓ minimizes the sum of vigilance costs and expected infection losses, solving

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

given others' random vigilance W one meets. In a *Bayes Nash equilibrium*, one solves optimization (1), and the random vigilance W is a best reply to the random loss L.

Since we will see that it suffices to make sense of the data, we assume that types mix uniformly and randomly in encounters. The opposite extreme when people sort by age formally reduces to Keppo et al. (2020), since matched individuals' losses coincide.

Static Analysis Suffices. Analyzing isolated periods is justified if (i) no one impacts future infection levels, as individuals are negligible (with a large population), and (ii) no one impacts their own future infection levels (re-infection happens). This logic ensures that no one solves a nontrivial dynamic optimization. So with falling re-infection chances, we would need to carry into a period a record of past infections.

3 Equilibrium Analysis

We consider everyone for whom the disease is dire enough — given its prevalence π and passing rate β: ($\bigstar$) *the support of losses* L *is in* (ℓ_0, ∞), where $\gamma\pi\ell_0\beta q = 1$. For when $\ell > \ell_0$, the marginal benefit of vigilance exceeds its marginal cost at $v = 0$.[5]

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 $V(\ell) > 0$ (and so the loss is $\ell > \ell_0$), the equilibrium FOC holds:

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

A solution of (2) is an optimum because $f'' > 0$, guaranteeing the SOC and a unique interior minimum (1). In other words, the vigilance function V is a best response to itself — namely, when facing the induced random vigilance $V(L)$. As every loss type ℓ optimizes, the vigilance function $V(\cdot)$ is a Bayes Nash equilibrium.

Let constant $C = \gamma\pi\beta qE[L^{-\gamma/(\gamma+1)}]$ reflect the passing rate, prevalence, and filter. We derive in §A the precise way that passing falls in equilibrium as π or ℓ or β rises:[6]

Theorem 1 (Equilibrium Passing) *The filter is* $f(V(\ell)) = C^{-\frac{\gamma}{2\gamma+1}}\ell^{-\gamma/(\gamma+1)}$ *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 greater losses above a threshold, namely, $\ell \geq \ell_0$, lead people to filter out more infections. While optimal vigilance $V(\ell)$ rises, increased losses are not fully displaced by vigilance. In other words, expected infection loss rises in ℓ, since $\ell f(V(\ell)) \propto \ell^{1/(\gamma+1)}$.

Hence, the vigilance game has the intuitive *strategic substitutes* property: If others' vigilance $W = V(L)$ increases (in the sense of first order stochastic dominance), then the marginal benefit of own vigilance is lower in (2). So the best reply $V(\ell)$ is lower.[7] This effort displacement is an essential characteristic of the Swiss Cheese Model.

Theorem 1 gives a log-linear formula for equilibrium incidence $I(\ell, \pi) = f(V(\ell))^2\beta\sigma\pi$.

Theorem 2 (Incidence) *In the unique Bayes Nash equilibrium, given* ($\bigstar$), *the* log incidence rate is

$$\log I(\ell, \pi) = B + \psi \log \ell + \varphi \log \pi \tag{3}$$

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

[5] In (2), we have $1 + f'(0)Ef(\tilde{V})\pi\ell_0\beta q > 1 - \gamma\pi\ell_0\beta q = 0$.

[6] So $f(V(\ell)) = (\gamma\pi\beta q\ell)^{-\frac{\gamma}{2\gamma+1}}$ as in Keppo et al. (2020) with one type $L \equiv \ell$, as $C = \gamma\pi\beta q\ell^{-\gamma/(\gamma+1)}$.

[7] Proof: 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.

This log-linear formula owes to the constant elasticity vigilance function.

Theorem 2 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 equation (3) are standard economic elasticities. As in the homogeneous agent BSI* model of Keppo 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 Keppo 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 relative to Keppo et al. (2020): ψ, the *incidence - infection loss elasticity*. Theorem 2 has three separately identifiable testable takeout messages, highlighting different aspects of the modeling:

1. Optimizing Behavior: $\psi < 0$ means that vigilance rises in the potential losses, while $\varphi < 1$ means that vigilance rises in prevalence, and thus infections rise less than the same percentage that prevalence rises. For a 1% higher disease loss makes one more vigilant, lowering infection rates by $|\psi|$%.

2. Matching externality: $-\psi < 1 - \varphi$, or avoidance effect of one's own losses is less than that of prevalence, since the latter impacts both meeting parties. So $|\psi| > 0$ because of one's own vigilance efforts, but the gap $|1 - \varphi| > 0$ reflects greater vigilance by two parties. For a 1% higher loss or a 1% prevalence rise each raise the expected infection loss 1%, given our optimization (1). But the latter commonly elicits more vigilance from *all matched parties*, and so depresses infections more.

This logic assumes everyone meets randomly-drawn types from the distribution. At the opposite extreme, we could venture that people sort by age. In this case, a greater loss always impacts two matched parties equally. We will next in §4 see that this is not the case, so that *infections are best understood as spread in random meetings rather than age-segregated encounters.*

3. Strategic substitutes property: $1 - \varphi < -2\psi$, or people react less to greater prevalence than twice one person reacts to the same percent loss increase. In other words, greater vigilance by one person displaces his partner's vigilance. This highlights how individuals understand that the strategic aspect of the pandemic, that others vigilance actions impact their own optimizations.

4 Data and Empirical Analysis

To apply our strategic model theory, we posit that people are rational, and minimize vigilance costs plus 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] In other words, the loss in optimization (1) is $\ell = \Delta\lambda$. So given expected utility and ignoring vigilance, our expected life loss in Theorem 2 equals βq times $\pi\ell = \pi\Delta\lambda$. Altogether, *people are equally harmed by an increase in the prevalence π and the same percentage rise in the death rate Δ — for either lifts expected losses $\pi\ell$ by the same percentage.* This explains why Keppo et al. (2020) can be repurposed for individual heterogeneity in losses λ.

4.1 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, 20–39, 40–49, 50–59, 50–69, 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

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

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

[10] Computations 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.

[11] We 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).

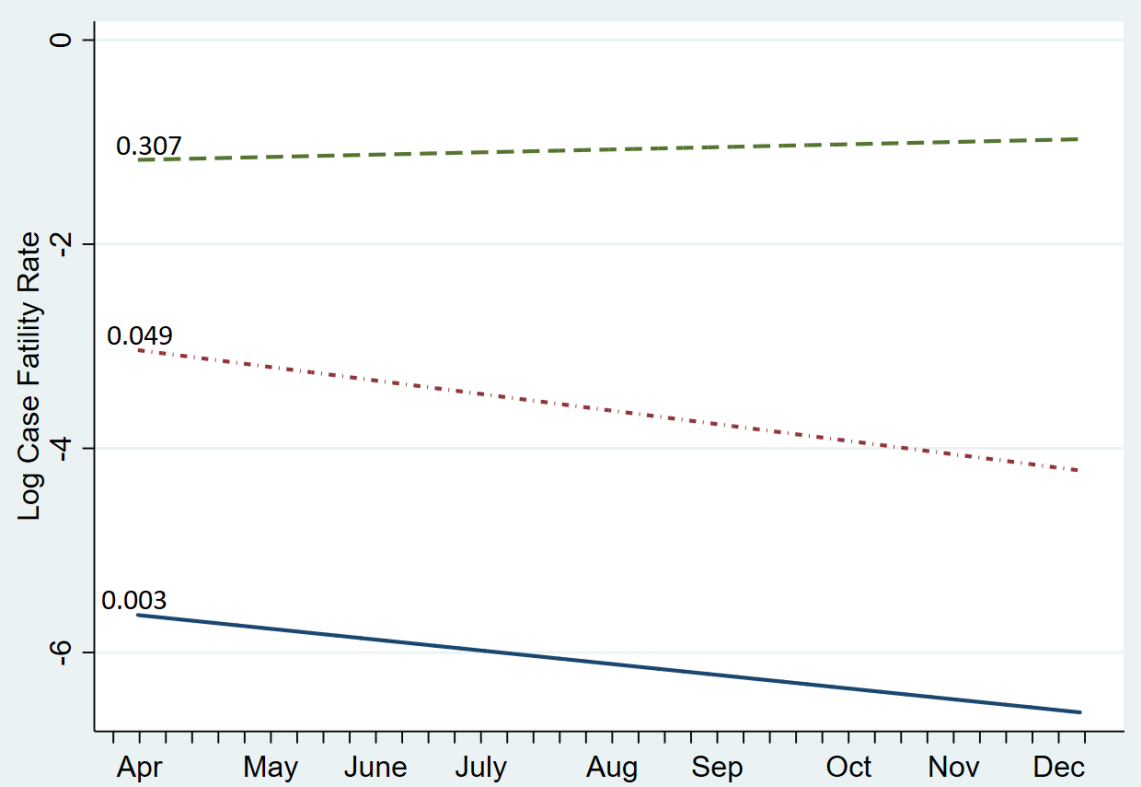

Figure 2: **The Estimated Log CFR with a Linear Time Trend.** We assume that individuals optimize on vigilance in response to their CFR. The age groups are 20-49 (solid) 50–69 (dot-dash), and 70+ (dash). The reported unlogged intercepts are roughly consistent with earlier noted CFRs. The slope estimates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

$i = 1, 2, 3$, corresponding to youngest, middle aged, and elderly.

$$\log(\text{CFR}_{it}) = a_i + b_i t + \epsilon_t$$

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]

4.2 Infection Rates

Finally, we turn to the COVID infection rates. To compute the infection rates by age, we use CDC data[14] 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 is reasonable to proxy the prevalence by the percent infected last week IR_{t-1}, and incidence by the rate this week. Then Theorem 2 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_i. The time fixed effects capture the driving

[12] Lacking such data for Massachusetts, we use France's age pyramid.

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

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

[15] Using 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.

[16] We 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.

CFR proxy $\longrightarrow$	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.

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 Keppo 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 2 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 Keppo 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.

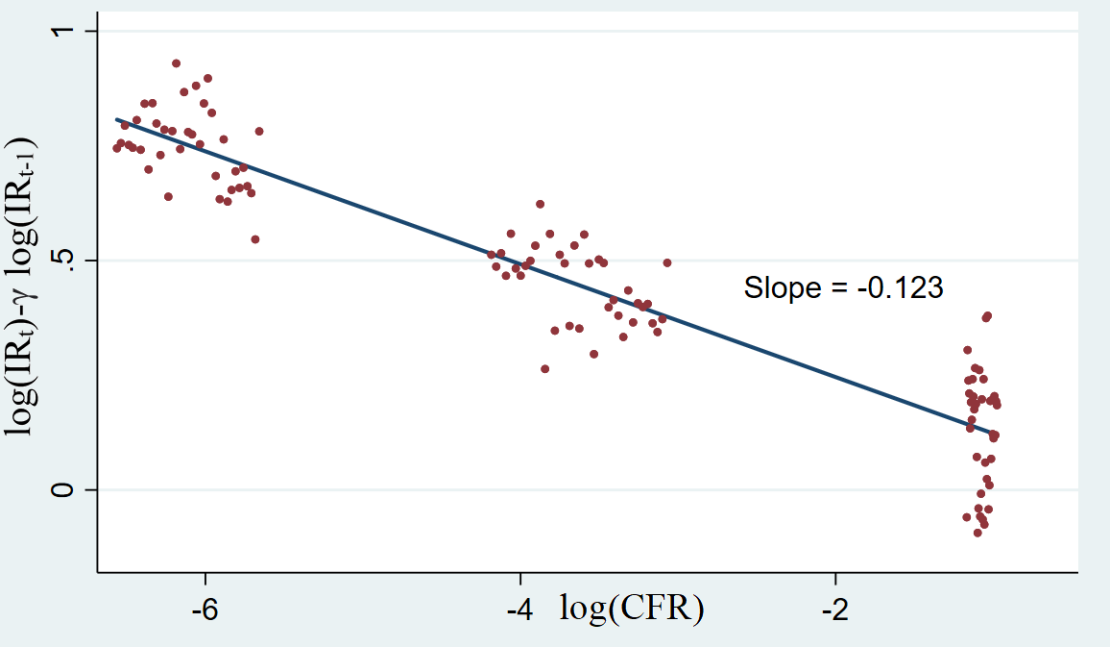

Figure 3: **The CFR Predicts the Infection Rate.** This is the scatter plot for the residuals of the regression of log infection rate on the CFR that are not explained by the last period infection rate, or by the age or time fixed effects.

5 Conclusion

Three events must happen a COVID death: exposure to the virus (prevalence), infection from it (passing), and death (or major adverse health outcome) conditional on infection. Keppo 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. In other words, *this paper offers a unified theory of how passing rates in contagions should respond to changes in the prevalence and in the death rate.*

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 Keppo 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]

As we emphasized earlier, this paper applies far afield from enriching contagion models with optimizing behavior. The Swiss Cheese model reflects any setting where accidents or failures reflect a systematic mistake by all people, such as auto accidents In these settings, more efforts by some parties can atone for less by others, and thus efforts are strategic substitutes. This is a fully solved Bayes Nash equilibrium of such settings, to capture heterogeneity in losses or mistake proclivities. For auto accidents, the variation among individuals will be in the cost of vigilance rather than loss. For worse drivers formally have higher vigilance costs. Young and very old drivers formally have a higher cost of vigilance. That is formally equivalent in our optimization to a lower accident loss, and our equilibrium will apply.

A Appendix: Omitted Proofs

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

$$f'(V(\ell)) + f''((V(\ell))V'(\ell)\ell = 0 \tag{4}$$

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

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

All that remains is to solve for the constant $c > 0$. Vigilance vanishes at the loss $\ell_0 > 0$ where $V(\ell_0) = c\ell_0^{1/(\gamma+1)} - 1 = 0$ implies $\ell_0 = c^{-(1+\gamma)}$. 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 $\ell \leq \ell_0$.

Since $\ell_0 = c^{-1-\gamma}$ by $V(\ell_0) = 0$ and equation (5), substituting $f'(V(t_0)) = f'(0) =$

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

$-\gamma$ into the FOC (2) yields:

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

So the constant c obeys $c^{2\gamma+1} = \gamma\pi\beta q E[L^{-\gamma/(\gamma+1)}]$. Finally, put $C = c^{2\gamma+1}$. □

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

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:

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

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 $\hat{I}(\ell,\pi)$. The 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.

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.

Keppo, Jussi, Marianna Kudlyak, Elena Quercioli, Lones Smith, and Andrea Wilson, "The Behavioral SIR Model, with Applications to the Swine Flu and COVID-19 Pandemics," 2020. mimeo.

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.