

Implications of Stochastic Transmission Rates for Managing Pandemic Risks
Harrison Hong, Neng Wang, and Jinqiang Yang
NBER Working Paper No. 27218
May 2020, Revised in September 2020
JEL No. G12,G32,Q5

ABSTRACT

We develop a model of pandemic risk management and firm valuation. We introduce aggregate transmission shocks into an epidemic model and link valuations to infections via an asset-pricing framework with vaccines. Infections lower earnings growth but firms can mitigate damages. We estimate a large reproduction number R_0 and transmission volatility for COVID-19. Using these estimates, we assess the accuracy of deterministic approximations based on R_0 . Our model generates predictions consistent with data: unexpected infection resurgence, non-monotonic mitigation policies, and higher price-to-earnings ratios during a pandemic. Valuations would be significantly lower absent mitigation and a high vaccine arrival rate.

Harrison Hong
Department of Economics
Columbia University
1022 International Affairs Building
Mail Code 3308
420 West 118th Street
New York, NY 10027
and NBER
hh2679@columbia.edu

Jinqiang Yang
Shanghai University of Finance
and Economics
Guoding Rd. 777
Shanghai, 200433
China
yang.jinqiang@mail.sufe.edu.cn

Neng Wang
Columbia Business School
3022 Broadway, Uris Hall 812
New York, NY 10027
and NBER
nw2128@columbia.edu

Implications of Stochastic Transmission Rates for Managing Pandemic Risks*

Harrison Hong[†] Neng Wang[‡] Jinqiang Yang[§]

September 13, 2020

Abstract

We develop a model of pandemic risk management and firm valuation. We introduce aggregate transmission shocks into an epidemic model and link valuations to infections via an asset-pricing framework with vaccines. Infections lower earnings growth but firms can mitigate damages. We estimate a large reproduction number \mathcal{R}_0 and transmission volatility for COVID-19. Using these estimates, we assess the accuracy of deterministic approximations based on \mathcal{R}_0 . Our model generates predictions consistent with data: unexpected infection resurgence, non-monotonic mitigation policies, and higher price-to-earnings ratios during a pandemic. Valuations would be significantly lower absent mitigation and a high vaccine arrival rate.

Keywords: COVID-19, pandemics, stochastic epidemic model, reproduction number, transmission volatility, vaccines, risk management, stock market valuation, mitigation, stochastic control

*We thank Ralph Koijen (Editor), two anonymous referees, Patrick Bolton, Ing-Haw Cheng, Marcin Kacperczyk, Jeffrey Kubik, Harry Mamaysky, Sen Pei, Bob Pindyck, Tom Sargent, Stijn Van Nieuwerburgh, Laura Veldkamp, and seminar participants at Columbia COVID-19 Virtual Symposium, Columbia Finance, and Imperial College London for helpful comments.

[†]Columbia University and NBER. E-mail: hh2679@columbia.edu

[‡]Columbia University and NBER. E-mail: neng.wang@columbia.edu

[§]Shanghai University of Finance and Economics. E-mail: yang.jinqiang@mail.sufe.edu.cn

1 Introduction

Deterministic epidemic models play a valuable role in the study of COVID-19 and its impact on the economy. On the public health front, epidemiologists recommend lockdown measures to keep \mathcal{R}_0 , defined as the expected number of secondary infections generated by a single (representative) infected individual in a fully susceptible population, below one (see, e.g., Ferguson et al. (2020)). On the economic front, economists model trade-offs of mitigation, where agents assume infection dynamics are deterministic (see, e.g., Atkeson (2020), Gourinchas (2020)).

Yet, both epidemiologists and economists recognize deterministic models are potentially crude approximations of stochastic epidemic dynamics. Aggregate transmission rate shocks due to environmental factors can play a large role in the evolution of infection dynamics (see Andersson and Britton (2012)). A case in point is the resurgence of COVID-19 in a number of countries during the summer of 2020, including countries with prudent mitigation such as South Korea. While epidemiologists recognize such stochasticity in fitting statistical models (see, e.g., Kucharski et al. (2020), Li et al. (2020)), it would be valuable to explicitly model how transmission volatility influence infection dynamics and optimal mitigation strategies.

This is especially important when one considers financial damage of the sort mentioned by the Federal Reserve Board Financial Stability Report (2020): “Asset prices remain vulnerable to significant price declines should the pandemic take an unexpected course...” Indeed, standard asset-pricing theory suggests that aggregate transmission shocks ought to play a critical role in asset prices through a risk premium channel.

Towards this end, we start with an extension of a widely-used deterministic epidemic model of COVID-19 (Kermack and McKendrick (1927)) featuring aggregate transmission-rate shocks that are intended to capture that viral contagiousness is unpredictable due to environmental factors.¹ Epidemic models of COVID-19 typically entertain multiple compartments in terms of tracking susceptible, infected, and resistant (including the recovered and dead). In order to transparently highlight the importance of transmission volatility and for

¹By environmental factors, epidemiologists broadly refer to weather, behavioral, cultural, and geographical factors.

tractability purposes, we focus on modeling just the infected population I_t , via a susceptible-infected-susceptible (SIS) as opposed to a susceptible-infected-recovered (SIR) setting.² For a number of economic and financial applications, the focus is typically on horizons of many years and the infected population is often the main state variable of interest since damages are likely to increase with infections. There is also no consensus at this point that COVID-19 infection confers long-lasting immunity.

Following Gray et.al. (2011), who have used this approach for other viruses, we model the COVID-19 aggregate transmission shocks via a stochastic transmission rate, $\tilde{\beta}_t$. This key input is modeled as a random variable with constant mean (predictable transmission captured by parameter β) and transmission shocks (mean zero but with volatility captured by parameter σ). The exit rate from the infected state back into the susceptible state is assumed to be a constant γ . We further introduce a stochastic vaccine arrival into our epidemic model as a jump with a Poisson arrival rate λ . When the vaccine arrives, we assume the pandemic is over and infections go to zero. The resulting dynamics of the fraction of infected then follows a four-parameter non-linear diffusion process.

We then model the financial impact of infection forecasts through the lens of a dynamic asset-pricing model. The unexpected arrival of COVID-19 directly affect earnings through three channels. First, there is a significant negative jump in earnings for firms in most industries, except for a few such as technology (Landier and Thesmar (2020)). Second, earnings growth is also potentially adversely affected by higher COVID-19 infections since higher infection rates mean their workers are sick and less productive. Third, stochastic transmission shocks also increase risk premia.

However, firms have access to a mitigation technology to reduce the drift of infections by paying both (flow) fixed and variable costs. Taking the stochastic discount factor used to price earnings as given, a representative firm optimally mitigates infections to reduce the damage on valuation via earnings, while taking into account the arrival of a vaccine. For instance, *MarketWatch* reported on May 2, 2020 that even a technology company like

²This SIS setting is useful for modeling viruses where recovery does not grant long-lasting immunity, which includes potentially many types of viruses.

Amazon will spend \$4 billion or more on COVID-19 mitigation responses such as testing for its workers, potentially wiping out the company's Q2 profit.

Reducing the disease spread incurs costs and hence lowers earnings in the short term but increases the expected earnings in the future by sufficient amounts so that it is optimal for the firm even from a value-maximizing perspective. As a result, even absent any health considerations, there are economic benefits to controlling COVID-19. Our model is a partial-equilibrium one and ignores externalities associated with mitigation. Our focus is justified by a number of studies including Andersen et.al. (2020) and Farboodi, Jarosch and Shimer (2020) that point to the importance of voluntary mitigation in terms of social distancing by households and firms even before government imposed lockdowns.³ Our model hence links risk management and valuation to epidemic data (infections, mitigation, reproduction numbers, transmission volatility, and vaccine arrival rate).

To relate our model to data, we first estimate our epidemic model by pooling COVID-19 case data from 16 countries (regions) that are at high risk during the period of January - February of 2020. These countries had among the most air travel connections to the city of Wuhan in China and have been the basis of the modeling of the early dynamics of COVID-19 before the onset of mitigation strategies. Given the noisiness and brief time series of the data and our goal of demonstrating the influence of shocks, we fit one model for all 16 countries.

Our estimate of β is 6.62 per month, which translates to an infected individual infecting one susceptible on average every five days ($\approx 30/6.62$.) Our estimate of monthly σ is 1.69, which translates to a standard deviation of plus or minus 1.69 infected individuals per month. The exit rate γ is equal to the inverse of the expected duration that an infected is sick and infective; it is typically not estimated based on aggregate data early in epidemics since there is a delay in when individuals leave the infected state. For our estimation of a population average, we simply use 14 days as the duration to infer the exit rate γ at $1/(14 \text{ days})$, which is 2.17 per month.⁴

³Externalities are also modeled in the COVID-19 context in Eichenbaum, Rebelo and Trabant (2020) and more generally for disasters in Hong, Wang and Yang (2020).

⁴This estimate ranges from around 7 days to 14 days at the individual level but with a fat-tail in terms of an infectious period across individuals. In empirical analysis, this parameter is typically assumed to follow

These estimates then imply that our (basic) reproduction number \mathcal{R}_0 using case data from January-February is 3.05 and the 90% confidence interval (CI) is (1.12, 6.52) based on the empirical distribution. Despite constraining one model for all countries (regions), our estimates are in line with leading studies of COVID-19.⁵

Along with these estimates that characterize the pre-mitigation COVID-19 process, we choose the remaining parameters to target asset-pricing moments and the magnitude of the decline in current earnings, which is nearly 40% for the typical firm (see Landier and Thesmar (2020)), and a post-mitigation reproduction number of around 1. We assume that the initial jump in earnings comes about equally from customers, who stop consuming when COVID-19 arrives and will not return until a vaccine arrives, and optimal mitigation on the firm. Our calibration assumes that the two channels contribute equally to the reduction transmission rates. We set the vaccine arrival rate λ to one per year, so that the expected pandemic duration is around one year based on surveys of vaccine experts and scientists.⁶

Our calibration generates a number of new insights. First, how well a deterministic model's infection forecasts approximates those of our stochastic model depends on vaccine arrival. Deterministic model infection forecasts tend to overshoot our model's conditional forecasts. The reasoning is understood by mathematical epidemiologists that introducing noise into the transmission process will lead to a dampening of stationary distribution of infections (Andersson and Britton (2012) and Brauer, Driessche, and Wu (2008)). Even if the reproduction number $\mathcal{R}_0 > 1$, the epidemic process might nonetheless die out due to the uncertainty of transmissions as opposed to in the deterministic setting. The further out the vaccine, the less good of an approximation is the deterministic model.

Second, our stochastic model yields rich optimal strategies beyond those from purely deterministic considerations. Because our estimated reproduction number is high, the optimal strategy always involves incurring fixed costs over a large range of infections rates, starting

an Erlang distribution (Kucharski et al. (2020)).

⁵Kucharski et al. (2020) estimate a reproduction number of 2.35 [95% CI 1.15-4.77], while Imai et al. (2020) estimate that it is 3.1 [95% CI 1.7-4.3].

⁶For such a timeline, see *McKinsey Report* (July 29, 2020) "On pins and needles: Will COVID-19 vaccines save the world".

at even a tiny fraction of the population being infected. However, transmission volatility generates an option value of waiting reflected by the optimal mitigation policy in the infection rate I_t on the intensive margin. Optimal mitigation policy can hence be non-monotonic in infection rates: initially increasing due to this option value of waiting and then declining due to standard congestion effects, since as infections rise, less of the population will then be susceptible. These implications match well with a standard epidemiological playbook that we have observed as we discuss below.

Third, however, the infection process cannot be perfectly controlled due to aggregate transmission shocks in contrast to deterministic models. Hence, mitigation at the intensive margin can then fluctuate with infection rates. But the optimally-mitigated COVID-19 infection process will tend to be pulled toward a constant reproduction number over time. These features match well the experience of a number of countries including the US that had to deal with unexpected resurgence of COVID-19 in early March, followed by a period where the reproduction number fell but then unexpectedly rose in the Summer of 2020.

Fourth, we show that the price-to-earnings (net of mitigation costs) can actually be higher during a pandemic than in a normal period consistent with data assuming there is optimal mitigation. The median annual price-to-earning ratio in the stock market has risen from around 19 before the pandemic to 24 during the pandemic. With both optimal mitigation and an expected vaccine arrival around one per year, mitigation while costly is temporary, earnings reduction is thus also temporary and moreover earnings is expected to discretely jump upward (due to both the elimination of mitigation costs and customers return once the pandemic is over. For these reasons, prices being the present value of all future earnings, fall much less than earnings do.

Fifth, we calculate a counterfactual for what would happen to stock prices absent any mitigation. The market would be down 15% relative to the optimal mitigation scenario. Our counterfactual can provide an account of asset price dynamics around COVID-19 — dramatic plunge in late February 2020 and equally dramatic rebound in March 2020. For instance, investors might have thought that society failed to control the virus initially (which corresponds to our counterfactual of no mitigation leading to lower stock market values) but

subsequently learned that they were going to (and hence the market rebounded to the equilibrium prices under optimal mitigation). Finally, we consider comparative statics exercises that further speak to observed asset price dynamics. Asset valuations are highly sensitive to the vaccine arrival rate, consistent at least with anecdotes of stock market response to news on vaccine developments.

Our paper proceeds as follows. We place our paper in related literatures in Section 2. We present our epidemic model in Section 3 and the valuation and risk management model in Section 4. We explore the early (pre-mitigation) dynamics of COVID-19 in Section 5. In Section 6, we highlight the role of stochastic transmission shocks, optimal mitigation and vaccine arrival for COVID-19's damage to stock valuations. In Section 7, we conduct counterfactual and comparative statics exercises. We conclude in Section 8.

2 Related Literatures

Our paper contributes to several literatures. In epidemiology, aggregate transmission shocks are used in epidemiological forecasting models to capture deviations of infections from deterministic projections (see Dureau et.al. (2013) for how parameter perturbation is employed). The analytical treatment of transmission volatility is fully articulated in Gray et al. (2011) who characterize the stationary distribution of our non-linear diffusion process absent a vaccine. Whereas $\mathcal{R}_0 > 1$ determines epidemic outbreak in a deterministic model, the analogous outbreak condition is $\bar{\mathcal{R}}_0 > 1$, where $\bar{\mathcal{R}}_0 = (\beta - \sigma^2/2)/\gamma$ in our stochastic model. Therefore, even at a reproduction number \mathcal{R}_0 above one, a sufficiently large value of σ^2 can cause $\bar{\mathcal{R}}_0 < 1$, which in turn implies that the epidemic becomes distinct in the limit. Zhao and Jiang (2014) extend the baseline model set-up in Gray et al. (2011) by allowing for a third compartment for the vaccinated. In their paper, a fraction of the population can become vaccinated. They explore stationary distributions in this setting.

Given that a goal of our analysis is to analyze the impact of an effective vaccine arrival on both disease spread and also valuation, we model vaccine via a Poisson jump process and explore the impact on conditional distributions, particularly how the deterministic model

approximations depend on this Poisson arrival rate, and also valuation. Hence, our model’s contribution to epidemiology is that it simultaneously accounts for a stochastic vaccine arrival, studies conditional distributions through Kolmogorov equations far from steady state, and derives the optimal mitigation strategy. Our optimally-mitigated stochastic SIS process is new to the epidemiology literature as there is limited work on stochastic control in epidemics.

There are other approaches to stochastic epidemics using Markov Chains. For instance, Allen and Burgin (2000) work with Markov Chain models with a discrete state space and one absorbing state so eventual disease extinction is guaranteed. That is, the stationary distribution is degenerate. Our model (with no vaccine) has a non-degenerate stationary distribution, which is better suited to explore implications of stochastic transmission shocks. Also our model features vaccine while theirs does not. Clancy (2005) also considers a Markov Chain SIS model set-up with no diffusion shocks where indirect transmissions (e.g., environmental bacteria and from animals to human) are possible. While this channel is relevant for some types of diseases (e.g., Zika, a mosquito borne flavivirus), it does not seem important for COVID-19.

In economics, recent theoretical models on controlling epidemics uses deterministic SIR models. The work closest to ours is Kruse and Strack (2020) where they show using a deterministic SIR model that the optimal solution is typically to act early unless herd immunity is within reach.⁷ Our contribution is to show how aggregate transmission shocks significantly influence optimal mitigation strategies in an SIS setting.

For tractability, we work with an SIS set-up rather than an SIR model as doing so yields an ODE rather than a PDE for the price-earnings ratio. While \mathcal{R}_0 is no longer a sufficient statistic in richer SIR, SEIR, and other even richer models with multiple compartments⁸ Our SIS model nonetheless captures first-order insights and mechanisms. As the mitigated I process tends to be low (near zero) most of the time in our calculations, the recovered

⁷Other models of optimal mitigation using a deterministic SIR model include among others Alvarez, Argente and Lippi (2020).

⁸A model with N compartments naturally calls for $N - 1$ state variables as the only restriction across the N compartments is that the population adds up to one (as a normalization).

fraction (in an SIR model) would also be close to zero. As a result, our approximation (which is effectively ignoring the recovered population) is likely to be sensible and hence our insights regarding \mathcal{R}_0 and deterministic approximations also have implications for SIR analysis.

In finance, our contribution is to link asset prices to underlying epidemiological data. Our model is consistent with Gormsen and Kojien (2020) who use a fundamentals-based asset-pricing model along with dividend futures to isolate a large impact of COVID-19 via the earnings growth channel. Hong, Kubik, Wang, Xiao and Yang (2020) combine our model with analyst forecasts to infer market expectations regarding the arrival rate of an effective vaccine that returns earnings to normal and to estimate the direct effect of infections on growth rates of earnings.

3 Stochastic Epidemic Model

In this section, for pedagogical purposes, we construct our stochastic model by starting with the classic Kermack and McKendrick (1927) model. Time is continuous and the horizon is infinite. We normalize the total population size to one and there is no birth nor death in the population. As a key motivation is to design a tractable and parsimonious model to conduct risk management applications, we only model two compartments (groups): infected and infectious (I) and susceptible (S) (or equivalently uninfected).⁹ Within each group, the population is homogeneous and well mixed. Let I_t and S_t denote the mass of the infected population and the susceptible at time t , respectively. As $I_t + S_t = 1$ at all t , we only need to keep track of the evolution for I_t , which is the single state variable in our model.

3.1 Deterministic SIS Model

Transmission rate in classic SIS setting. How does the disease get transmitted from an infected to a susceptible? The probability that an infectious individual meets a susceptible is

⁹In the epidemiology literature, there are various generalized formulations of these compartmental models. Widely used ones include SIR (susceptible, infected, recovered) and SEIR (susceptible, exposed, infected, and recovered) models. See Andersson and Britton (2012) and Brauer, Driessche and Wu (2008) for textbook treatments.

proportional to the product of their population mass, $I_t(1 - I_t)$, with an effective transmission rate, which we denote by β . Thus over the interval $[t, t + dt)$ the total number of new infections is

$$\beta I_t S_t dt = \beta I_t (1 - I_t) dt.$$

The infected recovers and becomes susceptible in our model. Let $\gamma > 0$ denote the rate at which an infected recovers. Hence, $1/\gamma$ is the duration for an infected to be infective. Subtracting the mass for the recovered $\gamma I_t dt$ over the interval $[t, t + dt)$ from the newly infected $\beta I_t (1 - I_t) dt$, we obtain the following process for dI_t , the net change of I_t :

$$dI_t = \beta I_t (1 - I_t) dt - \gamma I_t dt. \quad (1)$$

The solution of (1) satisfies the following logistic function:¹⁰

$$I_t = \left[\frac{\beta}{\beta - \gamma} (1 - e^{-(\beta - \gamma)t}) + \frac{1}{I_0} e^{-(\beta - \gamma)t} \right]^{-1}. \quad (2)$$

Basic reproduction number \mathcal{R}_0 . The basic reproduction number, \mathcal{R}_0 is defined as the expected number of secondary infections generated by a single (representative) infected individual in a completely susceptible population:

$$\mathcal{R}_0 = \frac{\beta}{\gamma}. \quad (3)$$

If $\mathcal{R}_0 \leq 1$ (when $\beta \leq \gamma$), the disease eventually is extinct, as (2) implies $\lim_{t \rightarrow \infty} I_t = 0$. If $\mathcal{R}_0 > 1$, the infected population I_t reaches the maximum level, $I_\infty = 1 - \mathcal{R}_0^{-1} > 0$ as $t \rightarrow \infty$, provided that $I_0 \neq 1 - \mathcal{R}_0^{-1}$.

We will use the term basic reproduction number and reproduction number interchangeably. The literature sometimes refers to the effective reproduction number at time t , which is the basic reproduction number multiplied by the susceptible mass. The effective reproduction number $\mathcal{R}_0 (1 - I_t)$ is time-varying in classic deterministic models.

¹⁰If $\beta = \gamma$, by applying the L'hospital's rule to (2), we obtain $I_t = \left(\beta t + \frac{1}{I_0} \right)^{-1}$.

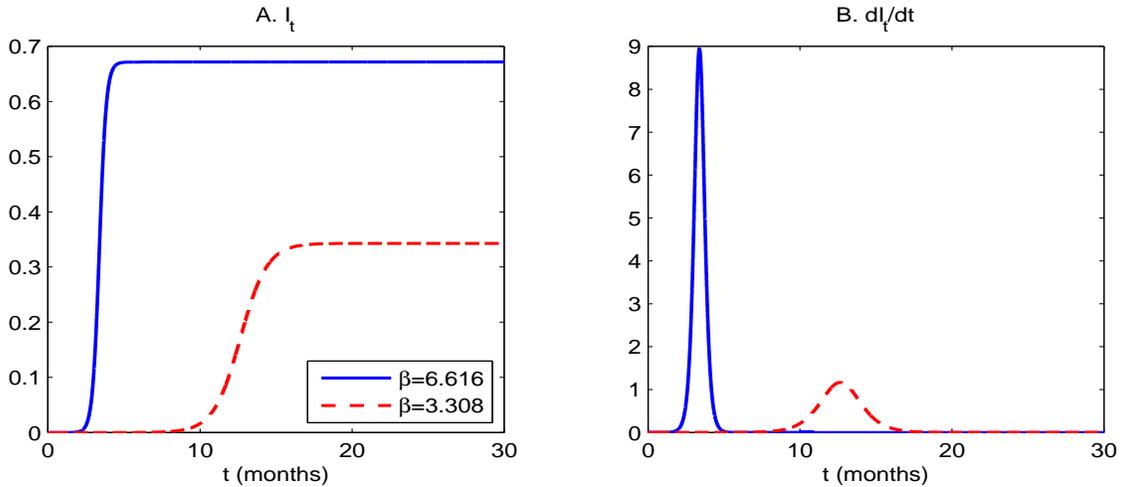


Figure 1: The infected fraction (I_t) and the net change (dI_t/dt) in a deterministic SIS model with $I_0 = 2 \times 10^{-7}$ based on the US data as of March 1st and $\gamma = 2.173$ per month.

Deterministic infection forecasts. In Figure 1, we plot the infected mass I_t at t in Panel A and the net change of the infected mass dI_t/dt in Panel B with the initial value of $I_0 = 66/(3.28 \times 10^8) = 2 \times 10^{-7}$ (as there were 66 infective individuals on March 1st in the US and the US population as of 2019 is 328 million.) The solid blue lines correspond to the solution for our deterministic case using our estimate of the transmission rate for COVID-19 that we discuss in Section 5.2. By reducing β by half from 6.616 to 3.308 per month, such as using economy-wide lockdowns, we lower the basic reproduction number \mathcal{R}_0 by half from 3.045 to 1.522 (unlike the three structural parameters, \mathcal{R}_0 is invariant to the time horizon we choose). As a result, the eventual infected fraction, I_∞ , decreases by half from 67.1% to 34.3% of the entire population.

Panel B captures the widely discussed flattening the curve argument (see, e.g., Atkeson (2020) and Gourinchas (2020)). Here, the curve refers to the net change of the infected population, dI_t/dt , as a function of time t . If the society successfully reduces β by half via social distancing and other interventions, this deterministic evolution curve is indeed significantly flattened and postponed. Specifically, this curve peaks at a bit over one year ($t = 12.657$ months) if $\beta = 3.308$ rather than at a bit over one quarter ($t = 3.384$ months). The curve of the net change, dI_t/dt , is substantially flattened.

Note the very sharp increase of I_t at the very early stage. This is because early on I_t is close to zero and we can thus effectively drop the $(1 - I_t)$ terms and approximate I_t as an exponential process: $dI_t \approx (\beta - \gamma)I_t dt$ with the approximate solution: $I_t \approx I_0 e^{(\beta - \gamma)t}$.

Obviously, exponential growth at a large rate $\beta - \gamma$ is incompatible with convergence of I_t to $I_\infty = 1 - \mathcal{R}_0^{-1}$ as $t \rightarrow \infty$. This is due to the dampening effect of I_t on its own growth. As the fraction of the infected increases, fewer are susceptible, which lowers dI_t/I_t . That is, the higher the level of I , the lower the infection growth rate dI_t/I_t .

3.2 Stochastic SIS Model

Aggregate transmission rate shock. Following Gray et al. (2011), a simple way to model stochastic transmission is to replace the constant rate β with a stochastic rate, which we denote by $\tilde{\beta}$. For expositional purposes, consider a discrete-time setting. The simplest choice for a stochastic $\tilde{\beta}$ is an independently and identically distributed (i.i.d.) random variable. Fix a small time increment Δ , we write

$$\tilde{\beta}_t \Delta = \beta \Delta + \sigma \sqrt{\Delta} \epsilon_t, \quad (4)$$

where both β and σ are constant parameters and ϵ_t is a mean-zero standard normal random variable.¹¹ Mapping (4) into our continuous-time formulation, we obtain

$$\tilde{\beta}_t dt = \beta dt + \sigma d\mathcal{Z}_t, \quad (5)$$

where \mathcal{Z}_t is a standard Brownian motion.

By using $\tilde{\beta} dt$ given in (5) to replace βdt in (1) and then combining drift and diffusion terms, we obtain the following stochastic differential equation (SDE) for I_t :

$$dI_t = [\beta(1 - I_t) - \gamma] I_t dt + \sigma I_t(1 - I_t) d\mathcal{Z}_t. \quad (6)$$

The drift term is the same as in the deterministic SIS model, while the diffusion term captures the uncertainty of the epidemiological evolution process. When no one is infected

¹¹By assuming ϵ_t is i.i.d., we make the transmission rate $\tilde{\beta}_t$ stochastic but without introducing an additional state variable for the transmission rate. We leave generalizations of our model to allow for a richer specification of $\tilde{\beta}_t$ for future work.

($I_t = 0$), the disease is extinct: $dI_t = 0$ as both drift and volatility terms in (6) are zero. If the entire population is infected ($I_t = 1$), the volatility has to be zero and the drift has to be negative so that the model is well posed.¹² Unlike $I_t = 0$, $I_t = 1$ is not an absorbing state as $\gamma > 0$.

Note that both the drift and volatility of the growth rate for the infected population, dI/I , depend on $(1 - I)$, the population of the susceptible. Specifically, the higher the level of I , the lower the drift (i.e., the expected infection growth rate) of dI_t/I_t . As the fraction of the infected increases, fewer are susceptible, which dampens the drift of dI/I .

To complete the description of our compartmental model, below we report the dynamics for the susceptible population S_t :

$$dS_t = (\gamma - \beta S_t) I_t dt - \sigma S_t I_t dZ_t. \quad (7)$$

Permanence of initial transmission shocks. The process for I_t given in (6) is not a Geometric Brownian motion (GBM) process widely used in Economics and Finance. But at a very early stage, I_t is close to zero; therefore, we can effectively drop the $(1 - I_t)$ terms in both drift and volatility functions and approximate I_t via a GBM process: $dI_t \approx (\beta - \gamma)I_t dt + \sigma I_t dZ_t$. That is, in the early stage, I_t evolves as

$$I_t \approx I_0 \exp \left[\left(\beta - \gamma - \frac{\sigma^2}{2} \right) t + \sigma Z_t \right] = I_0 e^{(\mathcal{R}_0 - 1)\gamma t} \exp \left(-\frac{\sigma^2}{2} t + \sigma Z_t \right). \quad (8)$$

Unlike the exponential growth approximation for I_t in the deterministic model, in our stochastic model, I_t is not only driven by \mathcal{R}_0 but also by the (exponential) martingale, the second exponential term in (8). This second term is equally important in driving the dynamics of I_t as the first (exponential) term involving \mathcal{R}_0 .

Because very few are infective early on, the change of I_t is highly idiosyncratic as the diffusion term dominates the drift term. A few negative shocks early on have outsized permanent effects on the evolution of I_t . On the other hand, if there are few such shocks early on, then the total infected population stays low for an extended period of time causing the

¹²If volatility is not zero or drift is positive at $I_t = 1$, the probability that I_t exceeds one is strictly positive, inconsistent with our model's assumption that the total population is normalized to one.

disease to be not that damaging. That is, in the very early stage, it is the sequence of realized values of $\tilde{\beta}$, not the expected transmission rate β used in the deterministic compartmental epidemic models, driving how fast the disease spreads.

Three-parameter non-linear diffusion process. We now have a three-parameter (β, γ , and σ) non-linear diffusion process. By applying Ito's Lemma to (6), we obtain:

$$d \ln I_t = q(I_t)dt + (1 - I_t)\sigma dZ_t, \quad (9)$$

where the drift for $\ln I_t$ is a quadratic function in I_t :

$$q(I) = \beta(1 - I) - \frac{\sigma^2}{2}(1 - I)^2 - \gamma. \quad (10)$$

Equations (9)-(10) are convenient to work with when we analyze the stationary distribution.

Stochastic Steady State (SS) and Stationary Distribution. Next we turn to the stochastic steady state and stationary distribution to gain some intuition for why \mathcal{R}_0 is an insufficient statistic for managing COVID-19 risks. The long-run distributional properties of the infected fraction I depend on all three parameters in a nonlinear way. Simply relying on \mathcal{R}_0 , which is ratio between the expected transmission rate β and exit rate γ can be quite misleading.

Unlike in the deterministic model, which generates a single number for I_t at any t , in order to fully capture the disease transmission dynamics, we next characterize the time-0 conditional distribution of I_t for all t . Let $f(I_t, t; I_0)$ denote the time-0 conditional density function for I_t , the infected mass at t given the initial infected mass I_0 .

Conditional distribution. The density function, $f(I, t)$, satisfies the following Kolmogorov forward equation:

$$0 = \frac{\partial f(I, t)}{\partial t} + \frac{\partial [(\beta(1 - I) - \gamma)If(I, t)]}{\partial I} - \frac{1}{2} \frac{\partial^2 [(\sigma I(1 - I))^2 f(I, t)]}{\partial I^2}. \quad (11)$$

The first term is the time effect on $f(I, t)$, the second term is the drift effect on $f(I, t)$, and the last term is the volatility effect on $f(I, t)$. In Section 5.3, we show how uncertainty substantially alters the transmission dynamics.

Next, we incorporate the impact of stochastic vaccine arrival.

3.3 Stochastic Vaccine Arrival

We assume that COVID-19 will disappear following a successful vaccine development.¹³ Specifically, we use the following SDE to model the evolution of I_t :

$$dI_t = [\beta(1 - I_{t-}) - \gamma] I_{t-} dt + \sigma I_{t-}(1 - I_{t-}) dZ_t - I_{t-} d\mathcal{J}_t. \quad (12)$$

We capture this vaccine arrival effect on I_t via the third term, where \mathcal{J}_t is a (pure) jump process with a constant arrival rate, which we denote by λ . When a vaccine is successfully developed, i.e., $d\mathcal{J}_t = 1$, the pandemic is extinct. We can generalize our model to allow for a multiple-stage vaccine development process with a gradual reduction of the infected population without losing much analytical tractability.¹⁴

4 Managing COVID-19 Risks

In this section, we develop a parsimonious yet operational model to capture the impact of pandemic shocks on fundamentals-based valuation. We show how COVID-19 parameters β (equivalently \mathcal{R}_0) and σ together with asset-pricing specifications impact valuation.

In Section 4.1, we propose a valuation model before unanticipated pandemic arrival. In Section 4.2, we consider optimal mitigation after an unanticipated pandemic arrival in an asset-pricing framework. We explicitly allow for stochastic vaccine arrival in our analysis.

4.1 Valuation Before Unanticipated Pandemic Arrival

To ease our exposition and set up the basic apparatus into which we later incorporate COVID-19 shocks, we first introduce a simple asset-pricing model with no pandemic shocks, i.e., under normal business-as-usual environment or when $I_t = 0$.

¹³The assumption that vaccine takes effect immediately is clearly made for simplicity. In reality, it may take a while for the population to be vaccinated and not everyone will be vaccinated. We can generalize our vaccine model to allow for a (large) fraction of the population to be vaccinated leaving a (small) fraction still susceptible (at the cost of additional complexity.) But the core of our analysis will remain valid.

¹⁴For brevity, we leave this extension out.

We start by specifying the following process for the widely-used simple stochastic discount factor (SDF), $\widehat{\mathbb{M}}_t$, in the normal regime:

$$\frac{d\widehat{\mathbb{M}}_t}{\widehat{\mathbb{M}}_t} = -r dt - \eta^{\mathcal{B}} d\mathcal{B}_t, \quad (13)$$

where \mathcal{B}_t is the standard Brownian motion for the aggregate shock. Here, r is the risk-free rate and $\eta^{\mathcal{B}}$ is the market price of risk for the aggregate shock.¹⁵ For simplicity, let r and $\eta^{\mathcal{B}}$ be constant. Equation (13) implies a one-factor model where the factor can be the aggregate consumption growth shock as in a representative-agent general-equilibrium model of Lucas (1978), or the market portfolio return in the CAPM of Sharpe (1964), the option pricing model of Black and Scholes (1973), or the portfolio choice problem of Merton (1971). Here, $\eta^{\mathcal{B}}$ is positive as a positive shock $d\mathcal{B}_t$ to the aggregate consumption growth or market return is good news which lowers the investor's marginal utility or equivalently $\widehat{\mathbb{M}}_t$.

Earnings process. We assume that the earnings process in the normal regime, \widehat{Y}_t , follows:

$$\frac{d\widehat{Y}_t}{\widehat{Y}_t} = \widehat{g}dt + \rho\phi d\mathcal{B}_t + \sqrt{1 - \rho^2} \phi d\mathcal{W}_t, \quad (14)$$

where \mathcal{B}_t is the aggregate shock introduced in (13) and \mathcal{W}_t is the standard Brownian motion driving earnings idiosyncratic risk. By construction, \mathcal{B}_t and \mathcal{W}_t are orthogonal. In (14), \widehat{g} is the expected earnings growth (drift) and ϕ is the volatility of earnings growth, which includes the aggregate component $\rho\phi$ and the idiosyncratic component $\sqrt{1 - \rho^2} \phi$. That is, ρ is the correlation coefficient between the aggregate shock \mathcal{B}_t and the asset's earnings process. For simplicity, we let \widehat{g} , ϕ , and ρ all be constant.

Pricing formula. Under the assumption that investors price earnings without expecting the possibility of a pandemic arrival, the firm's value in the normal regime (pre- and post-pandemic) satisfies the following standard asset-pricing equation (Duffie, 2001):

$$\widehat{P}_t = \mathbb{E}_t \left(\int_t^\infty \frac{\widehat{\mathbb{M}}_s}{\widehat{\mathbb{M}}_t} \widehat{Y}_s ds \right). \quad (15)$$

¹⁵No arbitrage requires that the drift of \mathbb{M}_t is equal to the minus interest rate, $-r$.

In Appendix C, using (13) and (14) and solving (15), we show that the firm's value is proportional to its earnings, $\widehat{P}_t = \widehat{p}\widehat{Y}_t$, where the price-earnings ratio, \widehat{p} , is a constant:

$$\widehat{p} = \frac{1}{r + \rho\phi\eta^{\mathcal{B}} - \widehat{g}}. \quad (16)$$

Equation (16) is the well-known Gordon growth model where $(r + \rho\phi\eta^{\mathcal{B}})$ is the firm's constant cost of capital (discount rate) and \widehat{g} is the earnings growth rate. This firm earns a risk premium of $\rho\phi\eta^{\mathcal{B}}$, which is given by the the product of the market price of risk $\eta^{\mathcal{B}}$ and $\rho\phi$, the systematic volatility component of ϕ and consistent with the one implied by the widely used CAPM.

4.2 Optimal Mitigation and Vaccine Arrival

Next, we incorporate pandemic shocks into our valuation model and consider the effect of mitigation responses by both customers (consumers) and firms.

After COVID-19 arrived unexpectedly in the U.S. late February, 2020, a fraction of consumers voluntarily engaged in social distancing and took various other precautionary measures reducing the spread of the disease. These voluntary actions by consumers (customers) substantially cut the transmission rate but also lower corporate earnings. Let Y_t denote the earnings process during the pandemic regime. To capture this phenomenon, we assume that logarithmic earnings drop at the moment of COVID-19 arrival time t_0 by a stochastic fraction $n(I_{t_0})$. That is,

$$Y_{t_0} = e^{-n(I_{t_0})} Y_{t_0-}, \quad (17)$$

where Y_{t_0-} is the pre-pandemic arrival earnings and Y_{t_0} is the post-pandemic arrival earnings.

The transmission rate is also lowered due to people's change of behaviors, e.g., quarantine, work from home, and social distancing reduce the speed of disease spreading. Let

$$\beta_{t_0} = (1 - \psi)\beta_{t_0-}, \quad (18)$$

where β_{t_0-} is the disease's transmission rate β (when the disease is not contained at all, and in the absence of any behavioral response.) The parameter ψ measures the fraction of β reduction due to customers' voluntary behavioral adjustments.

Next, we discuss the earnings process in the pandemic regime after t_0 .

Earnings process. Before the vaccine arrival at τ , we assume that Y_t is given by:

$$\frac{dY_t}{Y_{t-}} = g(I_{t-})dt + v(I_{t-})dZ_t + \rho\phi dB_t + \sqrt{1-\rho^2}\phi dW_t + (e^{n(I_{t-})} - 1)d\mathcal{J}_t. \quad (19)$$

COVID-19 influences Y_t as follows. First, the infection shock dZ_t directly causes additional earnings growth volatility, $v(I)$. Second, the expected earnings growth rate (absent vaccine arrival) is changed to $g(I)$ from \hat{g} . Third, the stochastic arrival of vaccine ($d\mathcal{J}_t = 1$) at $t = \tau$ causes an instantaneous (logarithmic) change of earnings from the pre-jump level, Y_{t-} , to the post-jump level \hat{Y}_t at $t = \tau$. For expositional simplicity, we assume that the percentage of earnings upward jump at the moment of vaccine arrival τ is equal to the percentage of earnings downward decrease at the moment of pandemic arrival time t_0 . That is, $\ln(\hat{Y}_\tau/Y_{\tau-}) = n(I_{\tau-})$, where $n(\cdot)$ is the same as the $n(\cdot)$ function in equation (17). We set $n(0) = 0$ so that earnings is continuous ($\hat{Y}_\tau = Y_{\tau-}$), when $I_{\tau-} = 0$.

Now consider a counter-factual case that helps us understand the mechanism of the model. Suppose $\tau- = t_0$, which occurs if $\lambda \rightarrow \infty$. For this case, earnings is not impacted at all by the two jumps (unexpected pandemic and vaccine arrivals) as the pandemic becomes extinct in no time $Y_\tau = e^{n(I_{\tau-})}Y_{\tau-} = e^{n(I_{t_0})}Y_{t_0} = e^{n(I_{t_0})}e^{-n(I_{t_0})}Y_{t_0-} = Y_{t_0-}$.

Additionally, earnings is still subject to the business-as-usual aggregate shock dB_t and idiosyncratic shock dW_t with volatility $\rho\phi$ and $\sqrt{1-\rho^2}\phi$, respectively. All shocks are orthogonal to each other.¹⁶ To highlight the role of stochastic transmission shocks on both earnings and valuation in a simple way, we assume that parameters for the business-as-usual aggregate variables and idiosyncratic risks, do not change with the unexpected pandemic arrival. We can of course also allow the business-as-usual parameters to also change as well without technical difficulties, but leave these extensions out for brevity.

¹⁶Indeed, $[\mathcal{W}_t, \mathcal{B}_t, \mathcal{Z}_t]^\top$ is a 3×1 standard Brownian motion and is independent of the vaccine arrival process \mathcal{J}_t .

SDF in the pandemic regime. As COVID-19 is clearly an aggregate shock, it changes the equilibrium SDF. We model the SDF \mathbb{M}_t in the pandemic regime as follows:

$$\frac{d\mathbb{M}_t}{\mathbb{M}_{t-}} = -r dt - \eta^{\mathcal{Z}} d\mathcal{Z}_t - \eta^{\mathcal{B}} d\mathcal{B}_t - (1 - e^{\kappa}) (d\mathcal{J}_t - \lambda dt). \quad (20)$$

As a positive pandemic shock $d\mathcal{Z}_t$ (which increases I) is bad news for the aggregate economy, investors' marginal utility (the SDF \mathbb{M}_t) should increase with I_t , which means $\eta^{\mathcal{Z}} < 0$, in contrast to a positive $\eta^{\mathcal{B}}$ for the business-as-usual aggregate shock $d\mathcal{B}_t$. The last term captures the effect of stochastic vaccine arrival on the SDF \mathbb{M}_t and this jump term is a martingale under the physical measure (to be consistent with no arbitrage).¹⁷ Upon the successful vaccine development at $t = \tau$, i.e., $d\mathcal{J}_t = 1$, the SDF immediately changes discretely from $\mathbb{M}_{\tau-}$ by $\mathbb{M}_{\tau} = e^{\kappa} \mathbb{M}_{\tau-}$. As a vaccine arrival is good news for the aggregate economy, investors' marginal utility (SDF) should decrease after vaccine arrival, which implies that the market price of vaccine arrival risk is negative, i.e., $\kappa < 0$.

Optimal stochastic control. Let $\{X_t; t_0 < t < \tau\}$ denote the mitigation process, which reduces the transmission rate, during the pandemic regime. Recall our assumption that the transmission rate is immediately reduced from the level of β (absent any behavioral response) to β_{t_0} by ψ fraction of β (see equation (18)). By spending X_t , the firm further lowers the transmission rate from β_{t_0} to $\beta_{t_0} - h_t$, where $h_t \geq 0$. This additional reduction obtained by mitigation X_t captures the effects of corporate actions.

Let $x_t = X_t/Y_t$ denote the scaled mitigation. We assume

$$h_t = h(x_t), \quad (21)$$

where $h(x)$ is increasing and concave in x . The motivation for the homogeneity (in earnings Y_t) assumption underpinning equation (21) is that to cut the transmission rate by the same magnitude $h(x_t)$, one firm whose earnings is twice the size of another needs to spend twice as much to achieve the same levels of reduction of the transmission rate. This assumption is reasonable as the benefit scales up with earnings and also makes our analysis tractable.

¹⁷Again, note that the drift rate of \mathbb{M}_t is equal to the negative interest rate by no arbitrage (Duffie, 2001).

The evolution of I_t given in (12) with mitigation is:

$$dI_t = [(\beta_{t_0} - h(x_{t-}))(1 - I_{t-}) - \gamma] I_{t-} dt + \sigma I_{t-}(1 - I_{t-}) dZ_t - I_{t-} dJ_t. \quad (22)$$

Note that β_{t_0} is the constant transmission rate after customers respond to the pandemic by taking precautionary measures. That is, absent mitigation but with customers' response, the transmission rate is lowered from β_{t_0-} to β_{t_0} . The basic reproduction number is then

$$\mathcal{R}_0(x_{t-}) = \frac{\beta_{t_0} - h(x_{t-})}{\gamma}. \quad (23)$$

Additionally, we assume that mitigation is costly and lowers earnings by more than the level of mitigation X . For simplicity we assume that mitigation incurs a (flow) fixed cost that is proportional to Y , i.e., πY_t where $\pi > 0$ is a constant measuring the size of flow fixed costs.

Let Y_t^* denote the firm's earnings netting out of both fixed and variable costs. With the above assumptions, the net earnings is

$$Y_t^* = Y_t - (\pi Y_t + X_t) \mathbf{1}_{X_t > 0}, \quad (24)$$

where $\mathbf{1}_{X_t > 0}$ is an indicator function that equals one if mitigation is strictly positive ($X_t > 0$) and zero otherwise.

Next, we state the optimization problem. The firm chooses mitigation X to maximize the following risk-adjusted present value in the pandemic regime:

$$P_t = \max_X \mathbb{E}_t \left(\int_t^\tau \frac{\mathbb{M}_s}{\mathbb{M}_t} (Y_s - (\pi Y_s + X_s) \mathbf{1}_{X_s > 0}) ds + \frac{\mathbb{M}_{\tau-}}{\mathbb{M}_t} \int_\tau^\infty \frac{\widehat{\mathbb{M}}_s}{\widehat{\mathbb{M}}_\tau} \widehat{Y}_s ds \right), \quad (25)$$

where τ is the stochastic vaccine arrival time. Inside the expectation operator in equation (25), there are two terms contributing to the discounted stochastic value of earnings: the first term is the value before the arrival of vaccine at τ and the second term is the value of after τ . Note that there is no need to spend on mitigation after τ but anticipation of vaccine arrival at τ fundamentally impacts the agent's optimal mitigation before τ .

Let C_t denote the corresponding present value of mitigation costs:

$$C_t = \mathbb{E}_t \int_t^\tau \frac{\mathbb{M}_s}{\mathbb{M}_t} ((\pi Y_s + X_s) \mathbf{1}_{X_s > 0}) ds. \quad (26)$$

4.3 Solution

Valuation in the pandemic regime. Because the earnings process features geometric growth, the firm's value is proportional to its earnings Y_t :

$$P_t = P(Y_t, I_t) = p(I_t)Y_t, \quad (27)$$

where $p(I_t)$ is the equilibrium price-earnings ratio.

As mitigating the disease involves a fixed flow cost πY_t , the firm may choose not to mitigate when the benefit of doing so is sufficiently low. Conditional on choosing $X > 0$, the optimal scaled mitigation, $x = X/Y$, satisfies the following first-order condition (FOC):

$$1 = -h'(x(I))I(1 - I)p'(I). \quad (28)$$

The level of mitigation $X_t > 0$ at the margin lowers the infected population mass I_t by $h'(x_t)I_t(1 - I_t)/Y_t$. Since the marginal increase of value from a lower level of I is $-P_I(Y_t, I_t) = -p'(I_t)Y_t > 0$, the marginal benefit of choosing $X_t > 0$ is equal to $-P_I(Y_t, I_t) \times h'(x(I_t))I_t(1 - I_t)/Y_t = -h'(x(I_t))I_t(1 - I_t)p'(I_t)$, the right side of equation (28). Optimal mitigation is chosen so that the marginal benefit of mitigation is equal to the marginal cost (which is one), i.e., $x(I)$ satisfies equation (28).

In Appendix C, we obtain the following ODE for $p(I)$:

$$0 = \max_{x \geq 0} - \left[(r + \rho\phi\eta^{\mathcal{B}} + v(I)\eta^{\mathcal{Z}} + \lambda^{\mathbb{Q}}) - g(I) \right] p(I) + 1 - (\pi + x)\mathbf{1}_{x>0} + \lambda^{\mathbb{Q}}e^{n(I)}\widehat{p} \\ + \left[(\beta^{\mathbb{Q}} - h(x) + v(I)\sigma)(1 - I) - \gamma \right] Ip'(I) + \frac{(\sigma I(1 - I))^2}{2} p''(I), \quad (29)$$

where

$$\beta^{\mathbb{Q}} = \beta(1 - \psi) - \eta^{\mathcal{Z}}\sigma \quad (30)$$

is the risk-adjusted disease transmission rate (i.e., under the risk-neutral measure \mathbb{Q}) and $\lambda^{\mathbb{Q}}$ is the risk-adjusted vaccine arrival rate (under the risk-neutral measure \mathbb{Q}):

$$\lambda^{\mathbb{Q}} = \lambda e^{\kappa}. \quad (31)$$

We expect the risk-adjusted duration of the pandemic (under the risk-neutral measure \mathbb{Q}), $1/\lambda^{\mathbb{Q}}$, to be longer than $1/\lambda$, which is the expected duration of the pandemic (under the

physical measure \mathbb{P}). This is because vaccine arrival is good news for the aggregate economy, which means $\kappa < 0$, as we discussed early.

Next, we turn to boundary conditions. First, as the no-infection state is absorbing, the price-earnings ratio at $I = 0$ is equal to the value in the normal regime:

$$p(0) = \widehat{p}. \quad (32)$$

Second, consider the extreme (and counter-factual case) where everyone is infected, $I = 1$. In this case, naturally there is no need to spend on mitigation, $x = 0$. The ODE (49) is then simplified as

$$[(r + \rho\phi\eta^{\mathcal{B}} + v(1)\eta^{\mathcal{Z}} + \lambda^{\mathcal{Q}}) - g(1)] p(1) = 1 + \lambda^{\mathcal{Q}}e^{n(1)}\widehat{p} - \gamma p'(1). \quad (33)$$

This boundary condition ties $p'(I)$ with $p(I)$ at $I = 1$.

Next, we go back to the optimal mitigation choice x . As the effective transmission rate depends on the product of I and $S = 1 - I$, mitigation is more valuable in an interior region of I , *ceteris paribus*. Given the fixed cost πY_t of mitigation, it is optimal for the firm to mitigate only when I is neither too high nor too low.

The preceding reasoning implies that the solution features three mutually exclusive regions: two inaction regions and one mitigation region. Denote \underline{I} and \bar{I} as the endogenously determined cutoff levels of the infected population mass for the three regions. In the regions $0 \leq I < \underline{I}$ and $\bar{I} < I \leq 1$, we have $x = 0$. In the region $\underline{I} < I < \bar{I}$, $x > 0$. The ODE for $p(x)$ in the mitigation region, i.e. $\underline{I} < I < \bar{I}$, is:

$$\begin{aligned} [(r + \rho\phi\eta^{\mathcal{B}} + v(I)\eta^{\mathcal{Z}} + \lambda^{\mathcal{Q}}) - g(I)] p(I) &= 1 - (\pi + x(I)) + \lambda^{\mathcal{Q}}e^{n(I)}\widehat{p} + \frac{(\sigma I(1 - I))^2}{2} p''(I) \\ &+ [(\beta^{\mathcal{Q}} - h(x(I)) + v(I)\sigma)(1 - I) - \gamma] I p'(I), \end{aligned} \quad (34)$$

where $x(I)$ satisfies equation (28).

For the two inaction regions, i.e. $0 < I < \underline{I}$ and $\bar{I} < I < 1$, $p(I)$ satisfies:

$$\begin{aligned} [(r + \rho\phi\eta^{\mathcal{B}} + v(I)\eta^{\mathcal{Z}} + \lambda^{\mathcal{Q}}) - g(I)] p(I) &= 1 + \lambda^{\mathcal{Q}}e^{n(I)}\widehat{p} + [(\beta^{\mathcal{Q}} + v(I)\sigma)(1 - I) - \gamma] I p'(I) \\ &+ \frac{(\sigma I(1 - I))^2}{2} p''(I). \end{aligned} \quad (35)$$

Finally, the endogenously cutoff levels, \underline{I} and \bar{I} , defining the three regions, satisfy:

$$p(\underline{I}_-) = p(\underline{I}_+), \quad p(\bar{I}_-) = p(\bar{I}_+), \quad (36)$$

$$p'(\underline{I}_-) = p'(\underline{I}_+), \quad p'(\bar{I}_-) = p'(\bar{I}_+), \quad (37)$$

$$p''(\underline{I}_-) = p''(\underline{I}_+), \quad p''(\bar{I}_-) = p''(\bar{I}_+), \quad (38)$$

where \underline{I}_- and \underline{I}_+ are the left limit and right limit of \underline{I} , and where \bar{I}_- and \bar{I}_+ are the left limit and right limit of \bar{I} . That is, $p(I)$ is twice continuously differentiable across all the three regions over the entire support $[0, 1]$.

In summary, $p(I)$ and the optimal mitigation policy $x(I)$ satisfy ODE (34) and equation (28) in the mitigation region $I \in (\underline{I}, \bar{I})$, and $p(I)$ satisfies ODE (35) in the inaction regions $I \in (0, \underline{I})$ and $I \in (\bar{I}, 1)$, where $x(I) = 0$. Solving equations (34), (28), and (35) subject to the five sets of boundary conditions given in equation (32) at $I = 0$, (33) at $I = 1$, and (36), (37), and (38) at the endogenously chosen thresholds, \underline{I} and \bar{I} , we fully characterize the model's solution over the entire support $I \in [0, 1]$ including \underline{I} and \bar{I} .

PV of mitigation costs. Once we have solved our optimization problem, we can calculate the present value of mitigation costs, $C(Y, I)$. Let $c(I) = C(Y, I)/Y$. Taking optimal $x(I)$, \underline{I} , and \bar{I} as given, and using essentially the same reasoning to calculate $p(I)$, we obtain the following ODE for $c(I)$:

$$\begin{aligned} [(r + \rho\phi\eta^B + v(I)\eta^Z + \lambda^Q) - g(I)] c(I) &= (\pi + x(I)) + \frac{(\sigma I(1 - I))^2}{2} c''(I) \\ &+ [(\beta^Q - h(x(I)) + v(I)\sigma)(1 - I) - \gamma] I c'(I). \end{aligned} \quad (39)$$

Because \underline{I} and \bar{I} are given for the purpose of calculating $c(I)$, we only require $c(I)$ to be once continuously differentiable at \underline{I} and \bar{I} :

$$c(\underline{I}_-) = c(\underline{I}_+), \quad c(\bar{I}_-) = c(\bar{I}_+), \quad (40)$$

$$c'(\underline{I}_-) = c'(\underline{I}_+), \quad c'(\bar{I}_-) = c'(\bar{I}_+). \quad (41)$$

Finally, the following boundary conditions hold at $I = 0$ and $I = 1$:

$$c(0) = 0, \quad (42)$$

$$[(r + \rho\phi\eta^B + v(1)\eta^Z + \lambda^Q) - g(1)] c(1) = 1 - \gamma c'(1). \quad (43)$$

5 Early Dynamics of COVID-19

5.1 Data

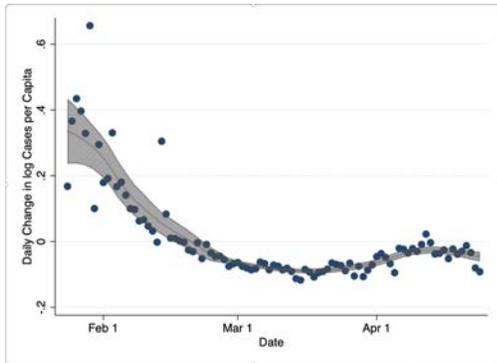
Our data on COVID-19 cases comes from COVID-19 Data Repository by Johns Hopkins available on github. The data keeps track of confirmed new cases, deaths, recoveries each day starting from January 22nd, 2020. The measure I_t in our model maps to the net number of outstanding infected cases at t , which is equal to the sum of the last period's I_{t-1} and the newly (reported) infected cases at t and subtracting deaths and recoveries, divided by the population of that country.

We follow leading epidemiological studies of COVID-19 and focus on China and countries (regions) that were at high risk due to air travel connected to China (Kucharski et.al. (2020)). There are a total of 16 countries in our sample. In Asia (Middle East), there are nine consisting of China, Japan, Malaysia, Singapore, South Korea, Taiwan (China), Thailand, United Arab Emirates and Vietnam. Among Western countries, these include Australia, Canada, France, Germany, Italy, United Kingdom, and United States.¹⁸

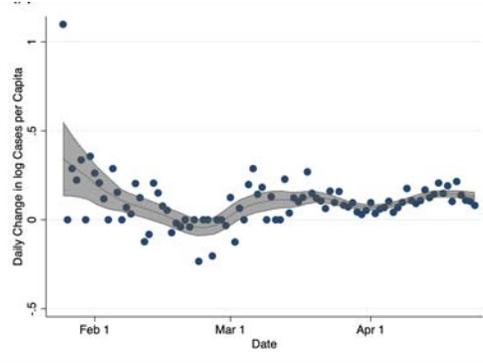
While all these countries have significant air travel connections to China, they did not experience the same infection path. This is consistent with our model that each country experienced idiosyncratic paths (realizations) of transmission shocks at early stages.

In Figure 2, we plot the logarithmic growth rate $d\ln(I_t)$ for four countries of interest, China, Singapore, US, and Italy. We can see that the epidemic curve of China reversed in the second half, while Singapore which had a fairly flat curve in the first half took off in the second half. The same is true for Italy and the US. Again, most of the countries in our sample only started mitigation responses later in the second half of the sample and it took time for them to have an effect.

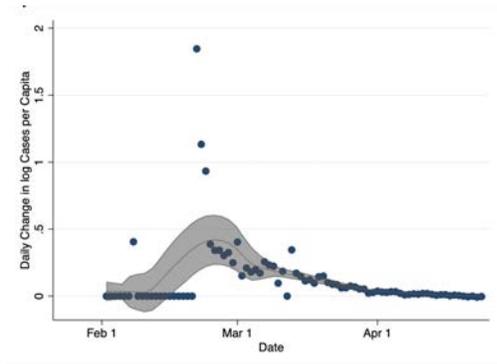
¹⁸Five of the original high-risk regions, Cambodia, India, Indonesia, Philippines and Russia, had no cases in January-February, so we exclude them from our analysis. These countries are also thought to be the most problematic in terms of underreporting of cases. See Manski and Monlinari (2020) for a discussion of biases in estimates of transmission rates associated with underreporting.



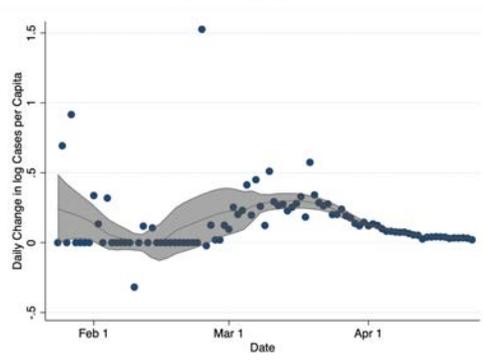
(a) China



(b) Singapore



(c) Italy



(d) US

Figure 2: Daily changes in $\ln I_t$, logarithmic infected population over time with shaded area being the 95-percent confidence intervals.

5.2 Estimation

We fit our model to the data for the January-February period. Given the noisiness and short time series of the data, we do not attempt to capture the potential heterogeneity in models across regions. Rather we think it is appropriate to fit one model by pooling the 16 regions. We pursue a robust estimation strategy as follows. For each region, we can estimate β and σ using the brief time series. But we use as our estimate the mean of the values across the 16 regions weighted by the number of daily observations in each region. For instance, China has more observations in the first sub-period and thus gets more weight in our estimate. We can then judge the sensibility of our estimates by comparing them to leading models of the early dynamics of COVID-19.

As we have pointed out a couple of times already, most mitigation responses only started

in March. Hence, we view our estimates as representative of the underlying epidemic process or early COVID-19 dynamics absent any mitigation responses.

Calibration of γ . Earlier epidemiological studies typically set γ by targeting the expected duration for an infected and infective individual to 14 days, which implies that the rate γ is equal to 1/14 per day, or $\gamma = 365/12/14 \approx 2.173$ per month. (Recall that in our convention, one period is one month.) Epidemiological studies typically view γ as highly predictable and relatively easy to estimate. They typically model this parameter as an Erlang distribution (Kucharski et.al. (2020)). By fixing γ , we leave out the impact of uncertainty of the exit rate on the disease spread.

Estimate of β . In Appendix A, we derive an OLS estimator for β given γ :

$$\hat{\beta} = \frac{1}{N-1} \sum_{i=0}^{N-2} \frac{I_{i+1} - 1 + \gamma}{1 - I_i}. \quad (44)$$

Table 1 reports the distribution of the estimate across regions. We use the mean estimate from the January-February sample (6.616 per month) as our baseline estimate with a 90% CI of (2.443, 14.168).

Table 1: The percentiles and moments for the monthly estimates of parameters β and σ^2 with implied \mathcal{R}_0 based on the data during the period of January-February 2020.

Estimates	Percentiles					Moments	
	5%	25%	50%	75%	95%	Mean	Std. Dev.
β	2.443	4.191	6.332	8.246	14.168	6.616	3.242
σ^2	0.718	1.138	1.436	4.791	8.857	2.851	2.537
\mathcal{R}_0	1.124	1.928	2.915	3.795	6.521	3.045	1.493

Estimate of σ^2 . In Appendix A, we show that the estimator for σ^2 is

$$\hat{\sigma}^2 = \frac{\sum_{i=0}^{N-2} (\ln I_{i+1} - \ln I_i)^2}{\sum_{i=0}^{N-2} (1 - I_i)^2}. \quad (45)$$

Table 1 reports the distribution of the estimate across regions. We use the mean estimate of σ^2 from the January-February sample (2.851 per month) as our baseline estimate with a 90% CI of (0.718, 8.857). The implied estimate of σ is then $1.689 = \sqrt{2.851}$ per month.

Estimate of the basic reproduction number \mathcal{R}_0 . Our estimate of the basic reproduction number \mathcal{R}_0 , shown in Table 1, is $3.045 = 6.616/2.173$ with a 90% confidence interval of (1.12, 6.52) based on data from the period of January-February. As we have mentioned in the Introduction, our estimates are in line with leading studies.

5.3 Vaccine Arrival Rates and Conditional Infection Distributions

We now focus on estimates and outcomes for the US, though our discussion equally applies to the aforementioned regions in our out-of-sample forecast analyses. We compare these conditional forecasts to the solution for the deterministic SIS model ($\sigma = 0$). In Table 2, we calculate the impact of vaccine arrival on the conditional mean and standard deviation of I_t for varying expected arrival rates.

We start at Panel A where there is no vaccine. We report the conditional means and variances for our stochastic SIS model with 1, 2, 4, 6, 8 weeks and 3, 4, 6, 9, 12, 24 months time horizons. Note that at the very early stage, e.g., from 1 week up to 6-8 weeks, the conditional mean forecast of I in our stochastic model is essentially the same as in the deterministic model — this is because the stochastic exponential growth approximation works well when I_t is very low. Starting from 3 months, this approximation no longer works. Deterministic model infection forecasts based on \mathcal{R}_0 overshoot our model’s conditional forecasts by a significant amount (0.104 for the deterministic model and 0.031 for our stochastic model). This is due to Jensen’s inequality. In around one to two years, the conditional distribution gets sufficiently close to the stationary distribution. While the conditional distribution results are new in our paper, the stationary distribution was characterized by Gray et al. (2011).

Furthermore, the conditional volatility is highly nonlinear and non-monotonic in the time horizon. For example, for the 3-month-ahead forecast, the monthly volatility of I_t (9.5%) is more than three times the mean (3.1%). Even with 6 months out, while the expected infected

Table 2: Effects of Vaccine Arrival Rate, λ , on Conditional Distributions of I_t . Other parameter values are: $\beta = 6.616$, $\gamma = 2.173$, and $\sigma = 1.689$ per month. The implied value of \mathcal{R}_0 is 3.045.

t	A. No Vaccine ($\lambda = 0$)			B. Six Months ($\lambda = \frac{2}{12}$)		
	Deterministic	Stochastic		Deterministic	Stochastic	
	I_t	$\mathbb{E}(I_t)$	$\sqrt{\text{Var}(I_t)}$	I_t	$\mathbb{E}(I_t)$	$\sqrt{\text{Var}(I_t)}$
1 wk	$5.6 * 10^{-7}$	$5.6 * 10^{-7}$	$5.4 * 10^{-7}$	$5.4 * 10^{-7}$	$5.4 * 10^{-7}$	$5.4 * 10^{-7}$
2 wk	$1.5 * 10^{-6}$	$1.5 * 10^{-6}$	$2.5 * 10^{-6}$	$1.4 * 10^{-6}$	$1.4 * 10^{-6}$	$2.5 * 10^{-6}$
4 wk	$1.2 * 10^{-5}$	$1.2 * 10^{-5}$	$4.4 * 10^{-5}$	$1.0 * 10^{-5}$	$1.0 * 10^{-5}$	$4.0 * 10^{-5}$
6 wk	$9.3 * 10^{-5}$	$9.2 * 10^{-5}$	$6.0 * 10^{-4}$	$7.3 * 10^{-5}$	$7.3 * 10^{-5}$	$5.4 * 10^{-4}$
8 wk	$7.1 * 10^{-4}$	$6.8 * 10^{-4}$	$5.2 * 10^{-3}$	$5.2 * 10^{-4}$	$5.0 * 10^{-4}$	$4.5 * 10^{-3}$
3 mo	0.104	0.031	0.095	0.061	0.019	0.076
4 mo	0.630	0.165	0.235	0.323	0.084	0.188
6 mo	0.671	0.519	0.251	0.248	0.191	0.294
9 mo	0.671	0.636	0.151	0.150	0.142	0.275
12 mo	0.671	0.639	0.146	0.091	0.086	0.225
24 mo	0.671	0.639	0.144	0.012	0.012	0.088
∞	0.671	0.639	0.144	0	0	0

t	C. One Year ($\lambda = \frac{1}{12}$)			D. Two Years ($\lambda = \frac{0.5}{12}$)		
	Deterministic	Stochastic		Deterministic	Stochastic	
	I_t	$\mathbb{E}(I_t)$	$\sqrt{\text{Var}(I_t)}$	I_t	$\mathbb{E}(I_t)$	$\sqrt{\text{Var}(I_t)}$
1 wk	$5.5 * 10^{-7}$	$5.5 * 10^{-7}$	$5.4 * 10^{-7}$	$5.5 * 10^{-7}$	$5.5 * 10^{-7}$	$5.4 * 10^{-7}$
2 wk	$1.5 * 10^{-6}$	$1.5 * 10^{-6}$	$2.5 * 10^{-6}$	$1.5 * 10^{-6}$	$1.5 * 10^{-6}$	$2.5 * 10^{-6}$
4 wk	$1.1 * 10^{-5}$	$1.1 * 10^{-5}$	$4.1 * 10^{-5}$	$1.2 * 10^{-5}$	$1.2 * 10^{-5}$	$4.2 * 10^{-5}$
6 wk	$8.2 * 10^{-5}$	$8.2 * 10^{-5}$	$5.7 * 10^{-4}$	$8.7 * 10^{-5}$	$8.7 * 10^{-5}$	$5.9 * 10^{-4}$
8 wk	$6.1 * 10^{-4}$	$5.8 * 10^{-4}$	$4.9 * 10^{-3}$	$6.6 * 10^{-4}$	$6.3 * 10^{-4}$	$5.1 * 10^{-3}$
3 mo	0.078	0.024	0.085	0.089	0.027	0.090
4 mo	0.450	0.117	0.213	0.531	0.138	0.225
6 mo	0.408	0.314	0.321	0.523	0.403	0.310
9 mo	0.318	0.300	0.334	0.462	0.437	0.321
12 mo	0.247	0.235	0.321	0.407	0.387	0.333
24 mo	0.091	0.086	0.225	0.247	0.235	0.320
∞	0	0	0	0	0	0

mass is 51.9% of the population, one standard-deviation is still 25.1%. The volatility declines once we go beyond 9 months out (Around 9 months, the volatility peaks at 0.321.) In other words, infection forecasts based on a deterministic model can be poor approximations for the conditional forecasts of our model for a large range of periods.

Importantly, the degree to which the deterministic model is a poor approximation depends on the vaccine arrival rate. In Panels B, C, and D, we consider the expected vaccine arrival time of 6 months, 1 year and 2 years. Consider Panel B where the vaccine is expected to arrive in six months. In this setting, infections at 4 months is 0.323 for the deterministic model and expected to be 0.084 for the stochastic model. In Panel A where the vaccine is not available, the corresponding numbers are 0.630 and 0.165.

Intuitively, when the initial fraction of the infected is low, the I_t process is approximately a Geometric Brownian motion with outsized drift β and volatility σ parameters. This means that shocks early on have permanent and very large effects. A vaccine that quickly arrives significantly truncates the duration of the pandemic substantially dampening the impact of shocks to this I process, thereby leading to the deterministic model being a better approximation.

More generally, while an eventual vaccine arrival will make the disease extinct, the conditional moments of I_t vary significantly with the expected arrival vaccine rate λ . Panel B shows that if the vaccine is expected to arrive soon (e.g., six months and hence $\lambda = 2/12$), the current conditional forecast of I_t for any horizon t is much lower than without vaccine arrival and peaks around 6 months at 19%. Panel C shows that if the vaccine is expected to arrive in one year and hence $\lambda = 1/12$, the current conditional forecast of I_t for any horizon t is much lower than the forecast without vaccine arrival (see Panel A with $\lambda = 0$). It peaks around 6 months at 31.4%, lower than the six-month-ahead forecast of 51.9% for the no-vaccine case, but higher than the six-month-ahead forecast of 19.1% for the more optimistic case with $\lambda = 2/12$.

Similarly, as we increase the expected waiting time for vaccine to two years (with a corresponding value of $\lambda = 0.5/12$ in Panel D) or no vaccine (Panel A), the expected infected fraction can reach as high as 44% and 64% around 9 months. These high numbers indicate

that waiting for vaccine to arrive will cause a very large fraction of the population to be infected absent behavioral responses even though in the long run COVID-19 goes extinct in our model due to the eventual arrival of vaccine by assumption.

6 Mitigation, Earnings, and Stock Market Valuation

We now present the solution of our model. We first discuss how we calibrate our model and then work through the four key predictions.

6.1 Calibration

Table 3 summarizes the parameter values of our calibration.

Earnings process. We set the annual earnings growth rate in normal times, \hat{g} , to 5% and the correlation coefficient between earnings growth and business-as-usual risk, ρ , to one.

Next, we specify the impact of the pandemic shock on the asset's earnings growth (drift) function $g(I_t)$ as follows:

$$g(I_t) = \hat{g} \left(1 - \zeta_1 I_t^{\zeta_2} \right), \quad (46)$$

where \hat{g} is the drift in normal times. The two parameters that govern the impact of infection rates on earnings growth are ζ_1 and ζ_2 . First, as $I_t = 0$ is an absorbing state, we set $g(0) = \hat{g}$ so that our pricing equation model is consistent with that under normal times. Second, earnings growth $g(I_t)$ is decreasing with I_t but at a slower rate as I_t increases implying $\zeta_1 > 0$ and $0 < \zeta_2 < 1$. While there is no historical data with which we can nail down these parameters, for our baseline we choose $\zeta_1 = 3$ and $\zeta_2 = 0.25$ so that $g(0.1) = -2/3 \times \hat{g}$ and $g(1) = -2 \times \hat{g}$. These are moderate long-run declines in growth rates absent mitigation. We set $v(I) = 0$ for simplicity. That is, infections will only affect the drift but not the volatility of earnings.

Finally, we specify the $n(I)$ function in the earnings process that appear at t_0 and τ as:

$$n(I) = \alpha_1 I^{\alpha_2}, \quad (47)$$

where $\alpha_1 > 0$ and $\alpha_2 > 0$. Recall that $I = 0$ is an absorbing state and $n(0) = 0$. We obtain $\alpha_1 = 18\%$ and $\alpha_2 = \frac{\ln(2/3)}{\ln(I_{t_0})}$ by targeting $n(I_{t_0}) = 12\%$, where $I_{t_0} = 2 \times 10^{-7}$, and $n(1) = 1.5 \times n(I_{t_0}) = 18\%$ at the moment of COVID-19 arrival time t_0 .

Asset-pricing parameters. We set the annual risk-free rate r at 4%, the annual stock market volatility σ_m at 20%, and the annual stock-market risk premium $(r_m - r)$ at 6%. The implied annual Sharpe ratio of the stock market $\eta^B = (r_m - r)/\sigma_m$, which is also the annual market price of business-as-usual risk in our CAPM model for the normal regime, is equal to $6\%/20\% = 30\%$. The implied asset's beta is $(\rho\phi/\sigma_m) \times (r_m - r) = \rho\eta^B\phi = 1 \times 30\% \times 20\% = 6\%$ and the cost of capital for this asset is equal to $4\% + 1 \times 6\% = 10\%$. As we set $\hat{g} = 5\%$, the price-earnings ratio in normal times, \hat{p} , is then equal to $1/(10\% - 5\%) = 20$, given by the Gordon growth model in equation (16).¹⁹

Next we set the pandemic asset pricing parameters. We set the market price of pandemic risk η^Z at -0.4 and the market price of vaccine arrival timing risk κ at -1 . As $v(I) = 0$, the only effect of η^Z is that for pricing purposes we need to use the risk-adjusted β^Q , which is different from β under the physical measure. Equation (30) implies the risk-adjusted transmission rate is larger than the real transmission rate for the pandemic (controlling for customers' voluntary precautionary response): $\beta^Q = \beta(1 - \psi) - \eta^Z\sigma > \beta(1 - \psi)$.

For the vaccine arrival rate λ , we set it at 1.1 per annum, consistent with optimistic assessment by scientists in the media. With $\kappa = -1$, the risk-adjusted vaccine arrival rate is $\lambda^Q/\lambda = 1.1 \times e^\kappa = 0.404$. That is, the risk-adjusted expected duration of the pandemic regime (under the risk-neutral measure \mathbb{Q}), $1/\lambda^Q$, is about two and half years, which is much longer than the expected duration of the pandemic regime under the physical measure, $1/\lambda = .9$ years.

Pre-mitigation COVID-19 dynamics, customers' response, and mitigation technology. We take the parameter values for COVID-19 dynamics absent mitigation from our

¹⁹For expositional simplicity, since we want to focus on the implications of vaccine arrival rates and stochastic transmission shocks, we have set the risk-free rate and business-as-usual risk to be the same across the normal and pandemic regimes. Of course, given the government interventions, these risks might have changed as well and one can adjust the calibration for different parameters in the pandemic regime.

estimates. For the period $t < t_0$, where t_0 is the arrival time of the COVID-19, the basic reproduction number (absent mitigation) is $\mathcal{R}_{t_0-} = \beta_{t_0-}/\gamma = 6.616/2.173 = 3.045$.

After the arrival of pandemic, customers' awareness and their voluntary cautionary measures automatically lower the speed at which the disease spreads, however, at the cost of lower earnings. Customers who stop consuming when the pandemic arrives and will not return until a vaccine arrives is assumed to lead to a decrease of earnings by a fraction of $1 - e^{-n(I_{t_0})} = 1 - e^{-0.12} = 11.3\%$, i.e., $n(I_{t_0}) = 12\%$ in equation (17).

The reduction of consumption also reduces the transmission rate of COVID-19. Even absent the firm's mitigation, the basic reproduction number decreases from \mathcal{R}_{t_0-} to $\mathcal{R}_{t_0} = \beta_{t_0}/\gamma = \mathcal{R}_{t_0-}(1 - \psi)$ due to customers' voluntary precautionary actions. By setting $\psi = 0.26$, customer awareness and precaution lower the basic reproduction number to $\mathcal{R}_{t_0} = 3.045 \times 0.74 = 2.25$ from $\mathcal{R}_{t_0-} = 3.045$.

Next, we parameterize the firm's mitigation technology. We assume that provided that the firm spends on mitigation, it incurs a flow fixed cost πY_t . We set $\pi = 5\%$. Conditional on choosing mitigation, the firm mitigation efficacy function takes the following form:

$$h_t = h(x_t) = \theta_1 x_t^{\theta_2}. \quad (48)$$

We set the parameter θ_1 to 20 per year (20/12 per month) and the curvature of firm mitigation efficacy θ_2 to 0.032.

6.2 Asset Pricing Absent Mitigation by Customers and Firms

Before analyzing the impact of mitigation (by customers and firms), we first summarize the asset valuation results absent any mitigation. This counterfactual exercise helps us better understand the value of mitigation on valuation.

Absent mitigation, the pricing ODE (34) for $p(I)$ can be simplified as:

$$\begin{aligned} [(r + \rho\phi\eta^B + v(I)\eta^Z + \lambda^Q) - g(I)] p(I) &= 1 + \lambda^Q \hat{p} + \frac{(\sigma I(1 - I))^2}{2} p''(I) \\ &+ [(\beta^Q + v(I)\sigma)(1 - I) - \gamma] I p'(I). \end{aligned} \quad (49)$$

As neither firms nor customers mitigate, $\mathcal{R}_0 = \beta/\gamma$ and $\beta^Q = \beta - \eta^Z\sigma$ as $\psi = 0$.

In Table 4, we focus on the first case where $\mathcal{R}_0 = 3.045$ absent any mitigation at all time and analyze the conditional dynamics of $p(I)$ as we vary the vaccine arrival rate λ .

Panel A shows that without the possibility of vaccine, the price-earnings ratios without and with stochastic transmission shocks are very close and they decrease from for the one-week ahead forecasts (5.53 and 5.64) to the long-run forecast (5.38 and 5.43), respectively. Absent vaccine and mitigation of any type, the stock market valuation reduction is catastrophic.

Moreover, absent a vaccine and mitigation, the valuation ratio declines with the forecast horizon due to a growth rate effect via a Gordon growth formula. We have assumed that growth rates are adversely affected by infection levels. Since infections are expected to rise over time, valuation ratios over time reflect increasingly lower growth rates. This result is easiest to see in the deterministic model—in the steady state, we have a high infection rate, negative growth and low valuation ratios.

Absent vaccine and mitigation of any type, the predictions for the deterministic and stochastic models are quite similar. We use the simpler deterministic model to shed lights on the mechanism. With no epidemic shocks, the I_t process converges to the steady state $I^{SS} = 1 - \mathcal{R}_0^{-1}$ as $t \rightarrow \infty$ and the equilibrium price-earnings ratio $p(I)$ at I^{SS} is given by

$$p(I^{SS}) = \frac{1}{r + \rho\phi\eta^{\mathcal{B}} - g(I^{SS})} = \frac{1}{r + \rho\phi\eta^{\mathcal{B}} - g_0(1 - \zeta_1 \cdot (I^{SS})^{\zeta_2})}, \quad (50)$$

where $g(I^{SS})$ is the long-term growth rate, r is the risk-free rate r , and $\rho\phi\eta^{\mathcal{B}}$ is the risk premium. This is a version of the Gordon growth model where the sum of the latter two terms is equal to the discount rate.

Quantitatively, with our baseline parameter values, I_t gets very close to $I^{SS} = 0.67$ at around 4.8 months (0.4 year) and the long-run growth rate is $g(I^{SS}) = -8.58\%$. Of course, this negative long-run growth rate is quite counterfactual. Later we show that vaccine arrival significantly alters the model's prediction. With $r = 4\%$ and $\rho\phi\eta^{\mathcal{B}} = 6\%$, we obtain $p(I^{SS}) = 5.384$ for the deterministic model. This highly counter-factual prediction hints the importance of the vaccine arrival and/or mitigation.

However, provided that an effective vaccine is expected to arrive and its arrival reverts

the economy to the normal, even absent any intervention, the price-earnings ratio becomes much closer to what we observe in the data. This is true even when investors expect an arrival of an effective vaccine around two years, which is on the longer end of forecasts, as shown in Panel D. With an expected vaccine arrival time of one year (Panel C) and stochastic transmission rates, the expected price-earnings ratio $p(I_t)$ is 15.8 (if t maps to one week), which is 21% lower than the pre-pandemic price-earnings ratio of 20. Intuitively, the forecasted $p(I)$ in the long run (steady state) is equal to the pre-pandemic value \hat{p} of 20. As such, with a vaccine, valuation ratios are rising over time since the vaccine arrival is expected to be closer the longer the forecast horizon.

Table 4 shows that even without any mitigation, whether the possibility of a vaccine arrival is priced in or not makes a fundamental difference on valuation. The pricing results with the possibility of vaccine arrivals are much closer to what we observe in reality. However, as we show later in this section, purely relying on vaccine without active mitigation by firms and/or customers, we are still unable to explain a number of other features of stock price dynamics. Nonetheless, the preceding counterfactual reported in Table 4 serves as a natural benchmark against which we calculate the value of optimal mitigation.

In Section 7, we conduct the counterfactual for the other case where firms do not mitigate but customers do. In this case, $\mathcal{R}_0 = \beta(1 - \psi)/\gamma$. We choose ψ to target $\mathcal{R}_0 = 2.25$.

6.3 Optimal Mitigation Policy

In Figure 3, we plot the optimal mitigation policy $x(I)$ as a function of $\ln(I)$. Because our estimated reproduction number is high, the benefit of mitigation is sufficiently high. The firm's optimal strategy involves incurring fixed costs over a large range of infection rates (i.e. $I \rightarrow 0$ and $\bar{I} = 0.947$). In our calibration, the firm mitigates even when infection rates are tiny. For example, even when $I = 2 \times 10^{-7}$, which means 66 infected in the US, the firm chooses $x(I) = 0.02$. Adding the fixed cost $\pi = 5\%$, the initial cost of mitigation is about 7% of the (gross) earnings.

Panel A shows that the optimal mitigation policy $x(I)$ is non-monotonic in infection

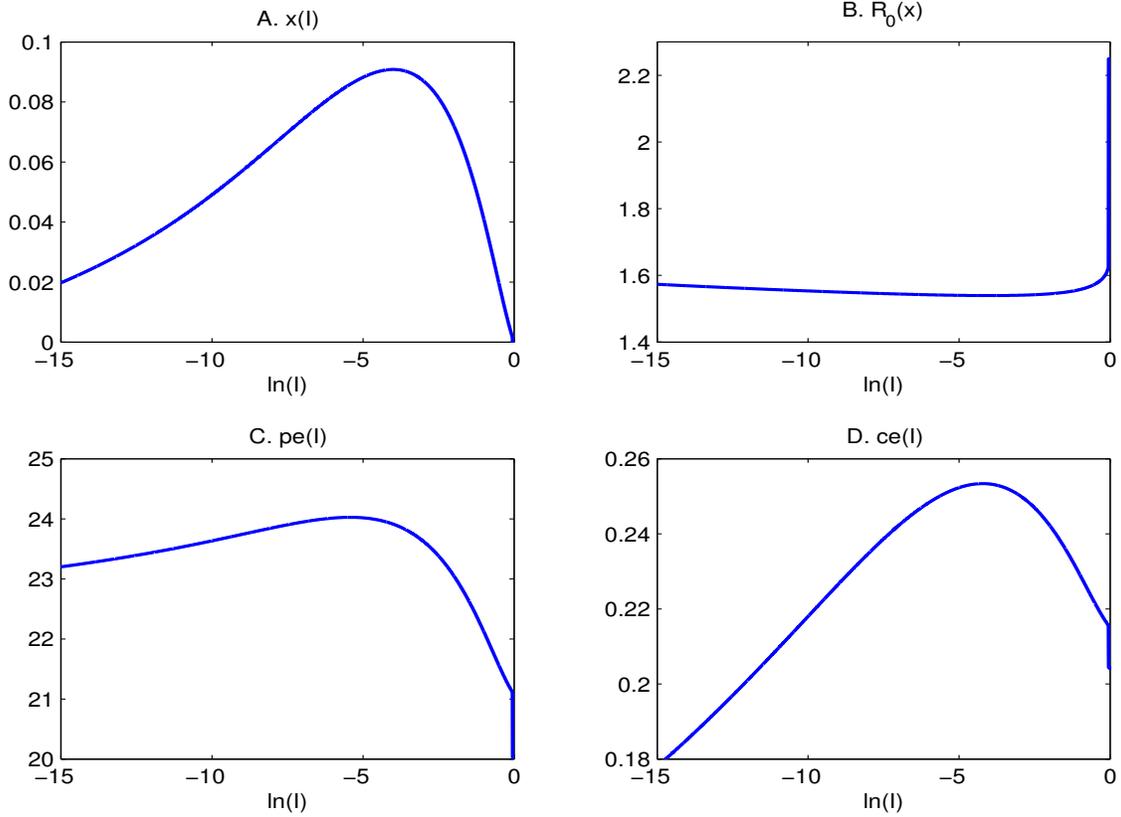


Figure 3: This figure plots the optimal mitigation policy, implied $\mathcal{R}_0(x)$, price-earnings ratio $pe(I)$ and scaled mitigation cost $ce(I)$. All parameter values are given in Table 3. The range for active mitigation is $I \in (\underline{I}, \bar{I})$, where $\underline{I} \rightarrow 0$ and $\bar{I} = 0.947$.

rates. It is initially increasing with I due to the option value of waiting. Transmission volatility implies there is a chance that the infections at very early stage will die out so it is optimal for the firm not to choose too high $x(I)$ for very low levels of I . Mitigation spending peaks when the (net) infected population reaches $I = 1.8\%$ (corresponding to $\ln(I) = -4$), at which point it then declines due to congestion effects since the susceptible population $(1 - I)$ decreases.

In Panel B, we see that mitigation ensures that the basic reproduction number $\mathcal{R}_0(x)$ stays at around 1.56. Only when $I_t \geq 0.947$, the firm no longer mitigates. As a result, $\mathcal{R}_0(x)$ goes back up to 2.25 which is the reproduction number when only customers mitigate and the firm loses around 11.3% of its (gross potential) earnings as a result.

In Panel C, we plot the price-to-earnings ratio $pe(I_t) = P_t/Y_t^*$ where Y_t^* is earnings net of mitigation costs. Because of optimal mitigation and a vaccine arrival rate of 1.1 per annum, the price-to-earnings ratio is actually higher during the pandemic than the price-to-earnings ratio \hat{p} in the normal period, which is 20 in our baseline specification. The reason is that earnings are net of mitigation costs which pulls down the denominator in $pe(I)$ but the price P_t is less affected by mitigation as vaccine arrival effectively shortens the duration of the pandemic making the COVID-19 shock transitory. Notice that $pe(I)$ first rises and then decreases with mitigation x , as mitigation $x(I)$ peaks for an interior value of I .

Finally, in Panel D, we also plot $ce(I_t) = C_t/Y_t^*$, the present value of mitigation costs scaled by earnings Y_t^* (net of mitigation costs) as a function of $\ln(I)$. The present value of these costs peaks also first increases and then decreases with infection rates. Notice moreover that even when $I \geq 0.947$, $c(I)$ remains positive and high since there is a chance that infections might fall to a lower value, at which point mitigation becomes valuable again.

6.4 Optimally-Mitigated Infection and Earnings Processes

To better understand our model's mechanism, we simulate a path of outcomes in Figure 4 with the initial value of $Y_{-0.5} = 1$ at $t = -0.5$. At $t = 0$, the pandemic unexpectedly arrives with an initial infection level of $I_0 = 2 \times 10^{-7}$. Panel A show a path of the optimally mitigated earnings process Y_t^* . It jumps down at $t = 0$ for two reasons. One is that some customers withdraw consumption and the other is the firm's optimal mitigation. As a result, earnings Y_t^* drops by around 20% at $t = 0$, remains relatively low during the pandemic ($0 < t < 1$) compared with the pre-pandemic period ($-0.5 < t < 0$), and then rebounds by a discrete amount when the vaccine arrives at $t = 1$ for two reasons, customers demand jump back up and the elimination of mitigation costs.

The infection process I_t cannot be perfectly controlled due to aggregate transmission shocks in contrast to deterministic models. While the COVID-19 infection process I is well managed most of the time, I_t still once in a while spikes up exceeding 1×10^{-3} , as we see in Panel B. Panels E and F show that mitigation at the intensive margin X_t and the present value of mitigation costs, C_t , fluctuates with infection rate I_t .

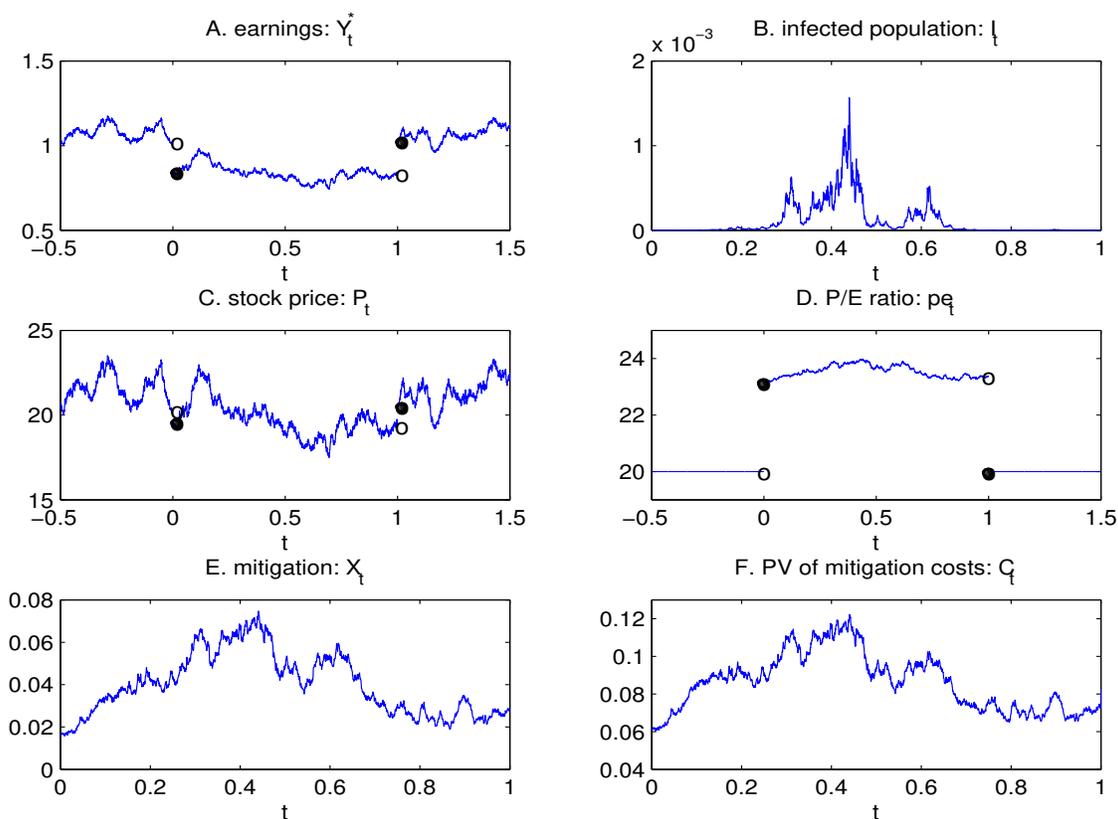


Figure 4: In this simulation, the pandemic unexpectedly arrived at $t = 0$ and the vaccine arrives at $t = 1$. Customers' precautionary responses lower the basic reproduction number from $\mathcal{R}_{0-} = \beta_{0-}/\gamma = 3.045$ (based on our pre-mitigation estimates) to $\mathcal{R}_0 = \beta_0/\gamma = 2.25$ absent the firm's mitigation. The simulation starts with $Y_{-0.5} = 1$ and the initial level of the infection rate is $I_0 = 2 \times 10^{-7}$ at $t = 0$. All parameter values are given in Table 3.

6.5 Stock Price and Price-to-Earnings Ratios

We next turn to stock price dynamics. Panel C of Figure 4 shows that the stock price P_t jumps downward at $t = 0$ when the pandemic unexpectedly arrives and upward at $t = 1$ when the vaccine arrives. In the interim, the higher infection realizations lead to lower asset prices. These features match well the experience of a number of countries including the US that had to deal with unexpected resurgence of COVID-19 in early March, followed by a period where the reproduction number fell but then unexpectedly rose in the Summer of 2020.

Moreover, Panel D shows more precisely that the price-to-earnings (net of mitigation

of costs), $pe_t = P_t/Y_t^*$, is actually higher during the pandemic than in the normal period consistent with data. The median price-to-earning ratio has risen from around 20 pre-pandemic to around 23.5 during the pandemic. With both optimal mitigation and a vaccine expected to arrive around one year, infections rarely get out of hand and the mitigation costs are temporary. Hence, prices do not fall much even as earnings Y_t^* are temporarily low.

7 Counterfactual and Comparative Statics

7.1 Counterfactual

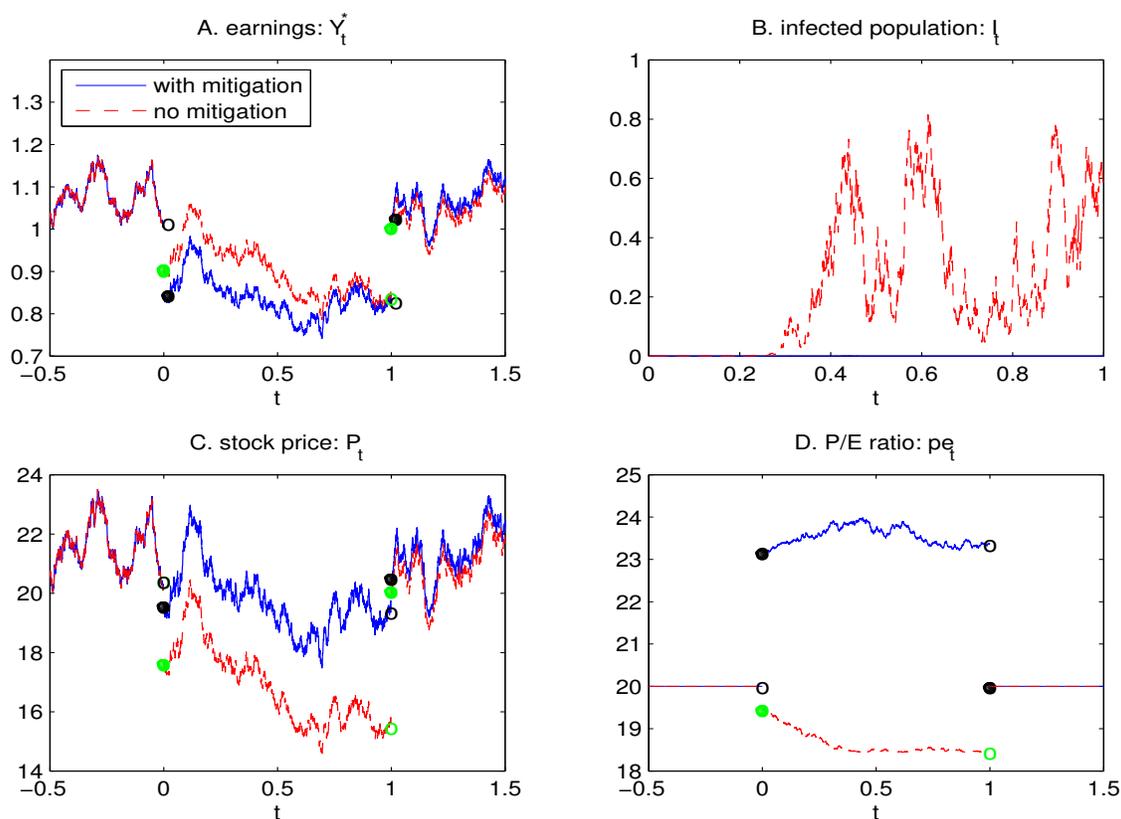


Figure 5: Counterfactual analysis: Comparing simulated paths for earnings (net of mitigation costs) Y_t^* , I_t , stock price P_t , and price-earnings ratio $pe_t = P_t/Y_t^*$ with and without mitigation. All parameter values are given in Table 3.

Finally, we demonstrate the effect of mitigation via a counterfactual in Figure 5 for what would happen to stock prices absent mitigation by the firm. This is the second case

described in Subsection 6.2, where customers' voluntary and precautionary measures causes the basic reproduction number \mathcal{R}_0 to drop from 3.045 to 2.25 once the public becomes informed about the pandemic at t_0 . The red lines depict the no-mitigation counterfactual. First, notice that at $t = 0$, earnings (net of mitigation costs), Y_t^* , falls less when the firm does not mitigate. But Y_t^* nonetheless still falls since some customers would withdraw purchases. Moreover, because the firm does not mitigate, the infection rate is much higher as we can see from Panel B. The earnings level with mitigation is initially lower than if the firm did not mitigate. However, earnings under the no-mitigation scenario deteriorates more than under the mitigation scenario. This of course is because infections get out of control with about 80% of the population infected at $t = 0.6$ in this simulation. This higher level of infections damages earnings growth. Once the vaccine arrives (in our simulation at $t = 1$), earnings actually jump more and end up being higher with mitigation than without mitigation. Because stock price P_t is proportional to earnings in the post-pandemic regime with the same coefficient given in equation (16), stock price also jumps more with mitigation than without at $t = 1$. In summary, the mitigation strategy outperforms the non-mitigation counterfactual because earnings and price are higher under the mitigation scenario when the vaccine arrives.

As a result, the stock market absent mitigation would be down by about 15% relative to the stock price under the optimal mitigation policy (Panel C). In contrast to the mitigation scenario, where the pe ratio hovers around 23.5, pe ratio is below 20 in the no-mitigation counterfactual (as seen in Panel D) due to multiple channels: damage to earnings growth and higher risk premia for stochastic transmissions. Hence, mitigation leads to less damage to stock markets.

Our model and this counterfactual in particular can give an account of asset price dynamics — dramatic plunge in the stock market in late February 2020 and equally dramatic rebound in late March 2020. An oft cited explanation for the rebound is that unexpected government intervention help alleviate elevated risk premia due to COVID-19 (Elenev, Landvoigt, and Van Nieuwerburgh (2020)). In our model, an important source of elevated risk premia and stock market volatility is aggregate transmission shocks.

But interventions by the Federal Reserve Bank in credit markets also coincided with a strong response by society to mitigate via social distancing. As such, our counterfactual also points to another explanation. Investors might have thought that society failed to control the virus initially (which corresponds to our counterfactual of no mitigation leading to lower stock market values) but subsequently learned that they were going to (and hence the market rebounded to the equilibrium prices under optimal mitigation).

7.2 Comparative Statics

Next, we consider several key comparative statics to gain further insights about our model.

Effect of transmission volatility σ . In Figure 6, we examine how our solution varies with transmission volatility σ . Our focus is on $x(I)$. First, Panel A shows that when the infected population is small, the higher the volatility, the lower the mitigation level $x(I)$. This is due to the option value of waiting at low levels of I . When I is low, the drift of I is positive and large. When transmission volatility σ is low, the dynamics of I is similar to a deterministic model. In the case with both low σ and low I , mitigating early and aggressively is optimal as it is very likely that I becomes large rather quickly (the drift of $\ln(I)$ is exploding). However, when I is low and σ is large, it is possible that the disease may die out on its own when there are only a few infected.²⁰ This explains why $x(I)$ decreases with σ when the infection rate I is low.

While the infected population I is large, the effect of σ on $x(I)$ reverses: the higher the volatility σ , the higher the level of mitigation $x(I)$. The intuition is as follows. When I is sufficiently high, the drift of I is negative due to the congestion effect (as the susceptible population $S = 1 - I$ is low) and mitigation decreases with I . Fixing I at a high value, a less volatile I process is more likely to take I to the left (as the negative drift effect of I plays a stronger role). As a result, the marginal benefit of mitigating when σ is low is smaller. This implies that for sufficiently high values of I , mitigation increases with σ .

²⁰Consider the scenario where a few infected all quarantine themselves. In this case, the disease becomes extinct.

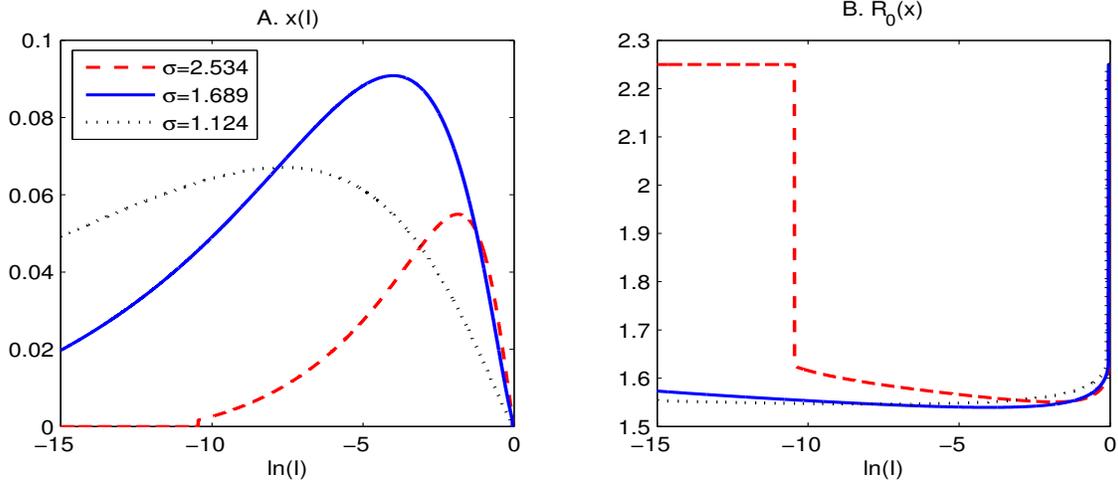


Figure 6: The effects of volatility σ on mitigation $x(I)$ and $\mathcal{R}_0(x)$. All parameter values are given in Table 3.

Panel B demonstrates the effects of volatility on $\mathcal{R}_0(x)$. As $\mathcal{R}_0(x)$ decreases with mitigation x , the effect of volatility σ on $\mathcal{R}_0(x)$ is the opposite of that on $x(I)$.

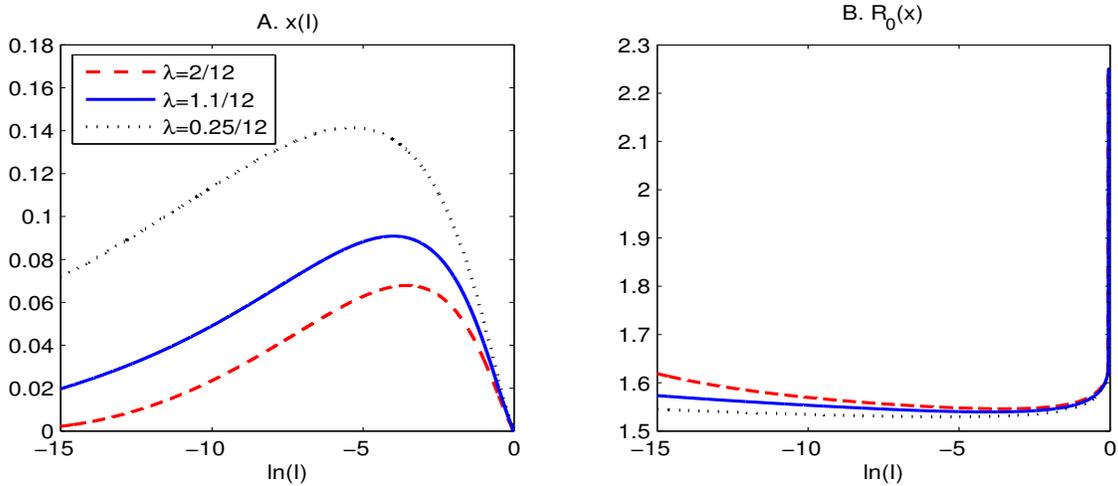


Figure 7: The effects of vaccine arrival rate λ on mitigation $x(I)$ and $\mathcal{R}_0(x)$. All parameter values are given in Table 3.

Effect of vaccine arrival rate λ . Figure 7 shows that the (scaled) mitigation level x_t decreases with the vaccine arrival rate λ (see Panel A). The intuition is as follows: The longer

the (expected) duration of the disease $1/\lambda$, the more valuable it is for the firm to spend on mitigation as shocks to I are highly persistent. Mitigating today lowers I which in turn lowers the future I in expectation. The longer the pandemic, the more valuable mitigation is and hence the higher the level of mitigation. The general non-monotonic pattern of $x(I)$ remains for a given λ .

Panel B demonstrates the effects of λ on $\mathcal{R}_0(x)$. As $\mathcal{R}_0(x)$ decreases with mitigation x , which decreases with λ , it is immediate to see that $\mathcal{R}_0(x)$ increases with λ . However, the quantitative effect is small. This suggests that the firm has a target \mathcal{R}_0 when choosing to mitigate. When the firm chooses not to mitigate, its $\mathcal{R}_0(x)$ shoots up as we see when I is near one.

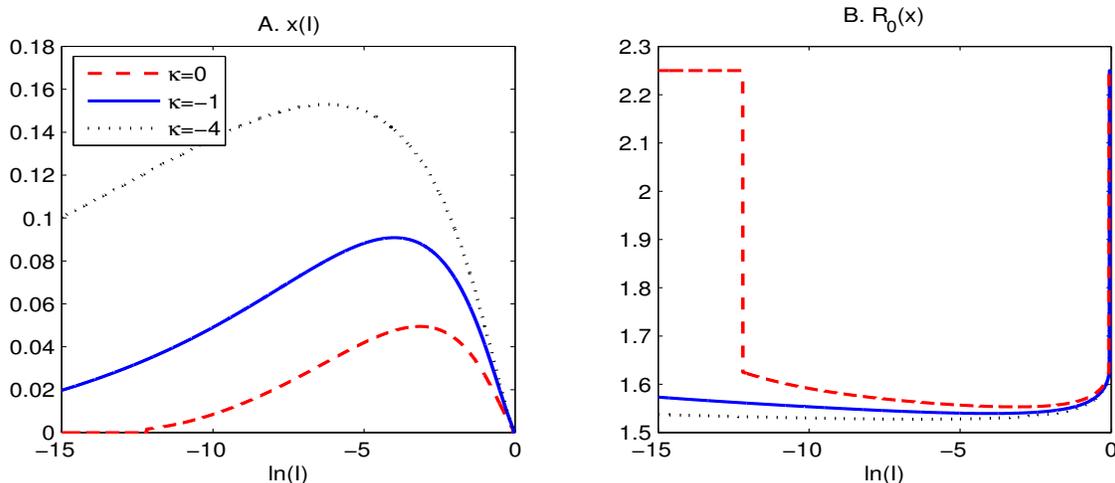


Figure 8: The effects of vaccine arrival timing risk premium κ on mitigation $x(I)$ and $\mathcal{R}_0(x)$. All parameter values are given in Table 3.

Effect of vaccine timing risk premium κ . Panel A of Figure 8 shows that the higher the vaccine arrival timing risk (a more negative value of κ), the more costly the pandemic shock is. As a result, the benefit of mitigating the disease increases with vaccine timing risk $|\kappa|$ and therefore mitigation x increases with $|\kappa|$ for a given level of I . The general non-monotonic pattern of $x(I)$ remains for a given κ .

Panel B demonstrates the effects of κ on $\mathcal{R}_0(x)$. As $\mathcal{R}_0(x)$ decreases with mitigation x ,

it is immediate to see that the higher the risk premium (a more negative κ), the lower the value of $\mathcal{R}_0(x)$. However, the quantitative effect is small. Again, this suggests that the firm has a target \mathcal{R}_0 when choosing to mitigate. When the firm chooses not to mitigate, its $\mathcal{R}_0(x)$ shoots up as we see when I is near one.

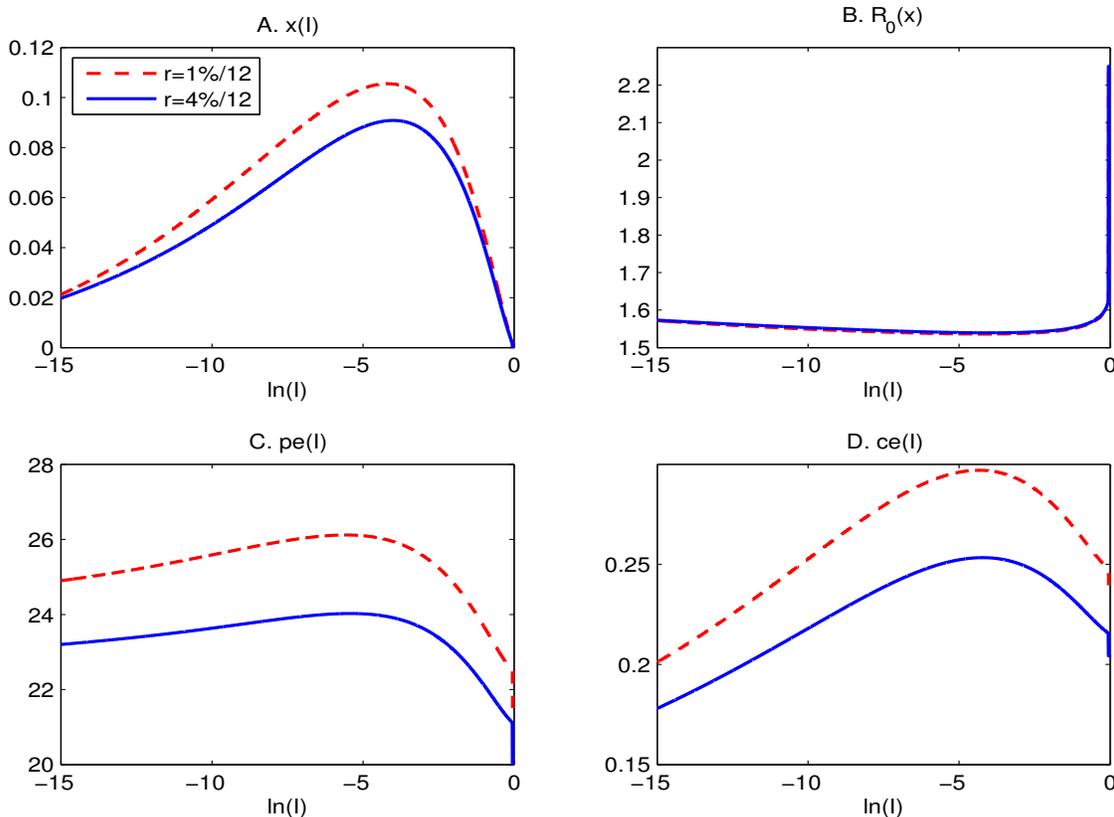


Figure 9: The effect of interest rate r on mitigation $x(I)$, $\mathcal{R}_0(x)$, price-earnings ratio $pe(I)$ and scaled mitigation cost $ce(I)$. All parameter values are given in Table 3.

Effect of lower interest rate r in the pandemic regime. In Figure 9, we show that as we decrease the (annualized) interest rate r (from 4% to 1%), the future earnings carry a greater weight for valuation, and as a result, mitigation $x(I)$ increases (Panel A). Panel B shows that the endogenous \mathcal{R}_0 is about the same for both levels of r . Panel C shows that the equilibrium price-earnings ratio $pe(I)$ as a result increases as r decreases. That is, as fiscal and monetary intervention measure were taken, $pe(I)$ naturally increases. Panel D

shows that the (scaled) present value of mitigation cost, $ce(I)$, also increases as the interest rate decreases, again for the same present-value argument: higher mitigation spending $x(I)$ discounted at a lower rate.

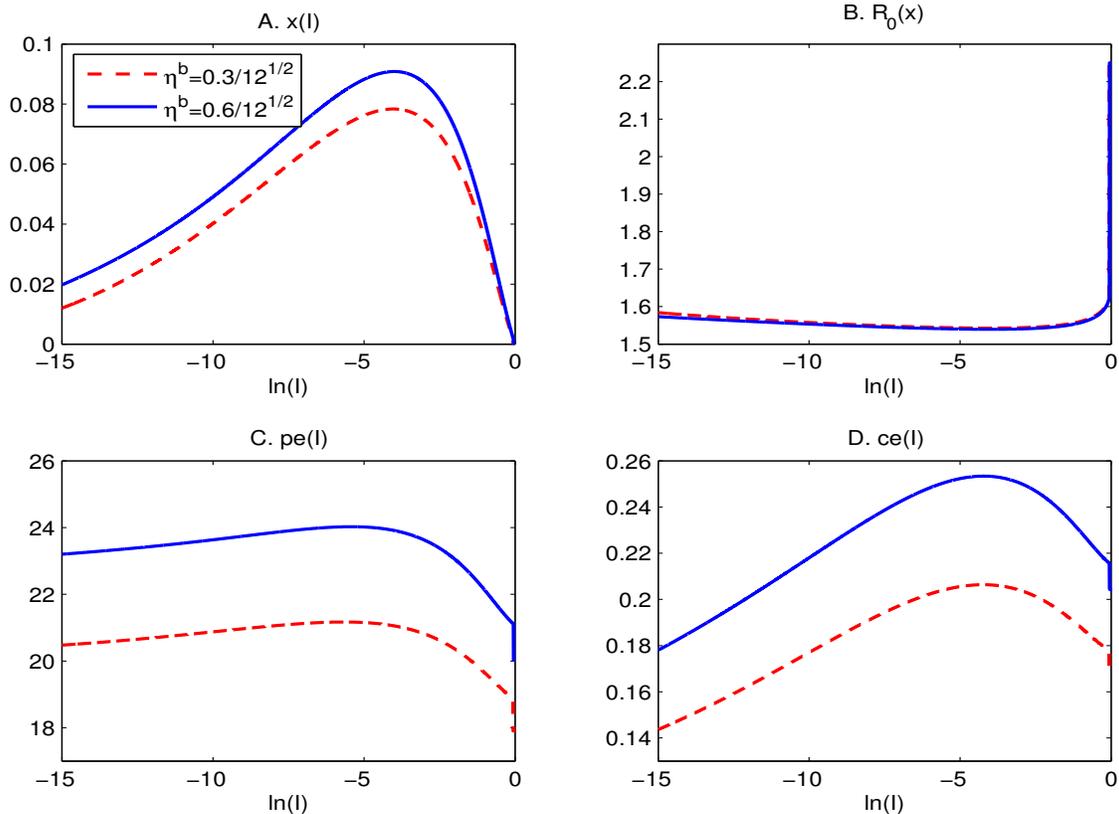


Figure 10: The effect of market price of risk η^B in the pandemic regime on mitigation $x(I)$, $\mathcal{R}_0(x)$, price-earnings ratio $pe(I)$ and scaled mitigation cost $ce(I)$. All parameter values are given in Table 3.

Effect of higher business-as-usual risk premium in the pandemic regime. In Figure 10, we show that as we increase the (annualized) market price of business as usual risk η^B (from 30% to 60%), the future earnings carry a smaller weight for valuation, and as a result, mitigation $x(I)$ decreases (Panel A). Panel B shows that the endogenous \mathcal{R}_0 is about the same for both levels of η^B . Panel C shows that the equilibrium price-earnings ratio $pe(I)$ as a result decreases as η^B increases. Panel D shows that the (scaled) present value of mitigation cost, $ce(I)$, also decreases as the market price of business as usual risk η^B increases, again

for the same present-value argument (a higher market price of risk discourages mitigation spending and lowers the present value, $ce(I)$.)

8 Conclusion

We propose a parsimonious epidemic model that highlights the importance of transmission-rate shocks due to unpredictable environmental factors and the prospect of stochastic vaccine arrival. The model is a four-parameter nonlinear diffusion process amenable for risk-management applications in areas such as economics and finance. We integrate the model into an asset-pricing framework that accounts for a potential vaccine and optimal mitigation so that we can quantify the financial damage of COVID-19 and relate this damage to epidemic data.

Our model has a number of implications for risk management including (1) the accuracy of deterministic model approximations depend on the vaccine arrival rate; (2) transmission volatility introduces an option-value-of-waiting consideration for the firm causing the optimal mitigation policy to be non-monotonic in infection rates; (3) a pandemic can lead to higher price-to-earnings ratios; and (4) stock market values would be significantly lower absent mitigation.

Appendices

A Estimation

Estimation of β . We use ordinary least squares (OLS) method to estimate the parameter β for a given value of γ . Discretizing I_t in (6) gives

$$I_{t+\Delta} = I_t + (\beta(1 - I_t) - \gamma)I_t\Delta + \sigma I_t(1 - I_t)\sqrt{\Delta} \epsilon_{t+\Delta}, \quad (\text{A.51})$$

where Δ is the time increment, $\epsilon_{t+\Delta}$ is a standard normal random variable, and

$$\frac{\left(\frac{I_{t+\Delta}}{I_t} - 1\right) - (\beta(1 - I_t) - \gamma)\Delta}{1 - I_t} = \sigma\sqrt{\Delta} \epsilon_{t+\Delta} \sim \mathcal{N}(0, \sigma^2\Delta). \quad (\text{A.52})$$

Let N denote the sample size. We choose an estimate of β to minimize the following:

$$\sum_{i=0}^{N-2} \left(\frac{\left(\frac{I_{(i+1)\Delta}}{I_{i\Delta}} - 1\right) - (\beta(1 - I_{i\Delta}) - \gamma)\Delta}{1 - I_{i\Delta}} \right)^2. \quad (\text{A.53})$$

The preceding equation implies that the estimate of β , $\hat{\beta}$, is given by

$$\hat{\beta} = \frac{1}{N-1} \sum_{i=0}^{N-2} \frac{\frac{I_{(i+1)\Delta}}{I_{i\Delta}} - 1 + \gamma\Delta}{1 - I_{i\Delta}}. \quad (\text{A.54})$$

And then by setting $\Delta = 1$ in (A.54) we obtain $\hat{\beta}$ given by (44). By setting $\Delta = 1$, the variance of $\hat{\beta}$ is given by

$$\text{Var}(\hat{\beta}) = \mathbb{E}(\hat{\beta} - \beta)^2 = \mathbb{E} \left(\frac{1}{N-1} \sum_{i=0}^{N-2} \frac{\frac{I_{i+1}}{I_i} - 1 + \gamma}{1 - I_i} - \beta \right)^2 = \frac{\sigma^2}{N-1}. \quad (\text{A.55})$$

The 95% confidence interval for $\hat{\beta}$ is therefore $\left(\hat{\beta} - 1.96 \frac{\sigma}{\sqrt{N-1}}, \hat{\beta} + 1.96 \frac{\sigma}{\sqrt{N-1}} \right)$.

Estimation of σ^2 . Equation (9) implies that the quadratic variation of $\ln I_t$, which we denote by $\langle \ln I_t, \ln I_t \rangle$, satisfies $d \langle \ln I_t, \ln I_t \rangle = (1 - I_t)^2 \sigma^2 dt$. Therefore, we have

$$\sigma^2 = \frac{\langle \ln I_t, \ln I_t \rangle}{\int_0^t (1 - I_s)^2 ds}. \quad (\text{A.56})$$

Discretizing the preceding equation, we obtain the following estimate of σ^2 :

$$\hat{\sigma}^2 = \frac{\sum_{i=0}^{N-2} (\ln I_{(i+1)\Delta} - \ln I_{i\Delta})^2}{\sum_{i=0}^{N-2} (1 - I_{i\Delta})^2 \Delta}. \quad (\text{A.57})$$

By setting Δ to one, we obtain (45).

B Deterministic SIS Model

Consider the case where $\beta \neq \gamma$. (The case with $\beta = \gamma$ is straightforward.) We have

$$\frac{dI_t}{dt} = \left[\frac{\beta}{\beta - \gamma} (e^{(\beta-\gamma)t/2} - e^{-(\beta-\gamma)t/2}) + \frac{1}{I_0} e^{-(\beta-\gamma)t/2} \right]^{-2} \left(\frac{1}{I_0} - \frac{\beta}{\beta - \gamma} \right) (\beta - \gamma) \quad (\text{B.58})$$

The second derivative of I_t is

$$\begin{aligned} \frac{d^2 I_t}{dt^2} = & -2 \left[\frac{\beta}{\beta - \gamma} (e^{(\beta-\gamma)t/2} - e^{-(\beta-\gamma)t/2}) + \frac{1}{I_0} e^{-(\beta-\gamma)t/2} \right]^{-3} \left(\frac{1}{I_0} - \frac{\beta}{\beta - \gamma} \right) (\beta - \gamma) \\ & \times \left[\frac{\beta}{2} e^{(\beta-\gamma)t/2} + \left(\frac{\beta}{2} - \frac{\beta - \gamma}{2} \frac{1}{I_0} \right) e^{-(\beta-\gamma)t/2} \right]. \end{aligned} \quad (\text{B.59})$$

Let t^* denote the time at which the peak of the net change dI_t/dt is reached, i.e., when $d^2 I_t/dt^2 = 0$. It is immediate to conclude that the curve dI_t/dt peaks at t^* where

$$t^* = \frac{1}{(\beta - \gamma)} \ln \left(\frac{\beta - \gamma}{\beta} \frac{1}{I_0} - 1 \right) = \frac{1}{(\beta - \gamma)} \ln \left(\left(1 - \frac{1}{\mathcal{R}_0} \right) \frac{1}{I_0} - 1 \right) \quad (\text{B.60})$$

C Derivation Details for Valuation

Valuation in the normal regime. In the normal regime no arbitrage implies that the drift under the physical measure for $\widehat{\mathbb{M}}_t \widehat{Y}_t dt + d(\widehat{\mathbb{M}}_t \widehat{P}_t)$ is zero (Duffie, 2001). By applying Ito's Lemma to this martingale, we obtain

$$\begin{aligned} 0 = & \widehat{\mathbb{M}}_t \widehat{Y}_t dt + \widehat{\mathbb{M}}_t \widehat{g} \widehat{Y}_t \widehat{P}'(Y_t) dt + \widehat{\mathbb{M}}_t \frac{\widehat{P}''(Y_t) \langle dY_t, dY_t \rangle}{2} + \widehat{P}(\widehat{Y}_t) (-r) \widehat{\mathbb{M}}_t dt \\ & + P'(\widehat{Y}_t) \langle d\widehat{Y}_t, d\widehat{\mathbb{M}}_t \rangle. \end{aligned} \quad (\text{C.61})$$

The above equation implies $\widehat{P}(\widehat{Y})$ satisfies the following pricing equation:

$$r\widehat{P}(\widehat{Y}) = \widehat{Y} + (\widehat{g} - \rho\phi\eta^B) \widehat{Y} \widehat{P}'(\widehat{Y}) + \frac{\phi^2 \widehat{Y}^2}{2} \widehat{P}''(\widehat{Y}). \quad (\text{C.62})$$

By solving the above ODE, we obtain $\widehat{P}(\widehat{Y}) = \widehat{p}\widehat{Y}$ where \widehat{p} is given by equation (16).

Valuation and mitigation in the pandemic regime. In the pandemic regime, no arbitrage implies that the drift under the physical measure for $\mathbb{M}_{t-} (Y_{t-} - (\pi Y_{t-} + X_{t-}) \mathbf{1}_{X_{t-} > 0}) dt + d(\mathbb{M}_t P_t)$ is zero. The pricing function depends on both Y_t and I_t . Hence, we write $P_t =$

$P(Y_t, I_t)$. By applying Ito's Lemma to this martingale, we obtain

$$\begin{aligned}
0 = & \mathbb{M}_{t-} (Y_{t-} - (\pi Y_{t-} + X_{t-}) \mathbf{1}_{X_{t-} > 0}) dt \\
& + \mathbb{M}_{t-} [((\beta(1 - \psi) - h(x_t))(1 - I_{t-}) - \gamma) P_I(Y_{t-}, I_{t-}) + g(I_{t-}) Y_{t-} P_Y(Y_{t-}, I_{t-})] dt \\
& + \mathbb{M}_{t-} \left[\frac{P_{II}(Y_{t-}, I_{t-}) < dI_t, dI_t >}{2} + P_{IY}(Y_{t-}, I_{t-}) < dI_t, dY_t > + \frac{P_{YY}(Y_{t-}, I_{t-}) < dY_t, dY_t >}{2} \right] \\
& + P(Y_{t-}, I_{t-}) (-r + (1 - e^\kappa) \lambda) \mathbb{M}_{t-} dt + P_I(Y_{t-}, I_{t-}) < dI_t, d\mathbb{M}_t > + P_Y(Y_{t-}, I_{t-}) < dY_t, d\mathbb{M}_t > \\
& + \mathbb{M}_{t-} [e^\kappa P(e^{n(I_{t-})} Y_{t-}, 0) - P(Y_{t-}, I_{t-})] \lambda dt.
\end{aligned}$$

Simplifying the above equation, we obtain the following pricing equation:

$$\begin{aligned}
rP(Y, I) = & \max_{x \geq 0} Y - Y(\pi + x) \mathbf{1}_{x > 0} + [g(I) - v(I) \eta^Z - \rho \phi \eta^B] Y P_Y(Y, I) \\
& + [(\beta(1 - \psi) - h(x) - \eta^Z \sigma) (1 - I) - \gamma] I P_I(Y, I) \\
& + \frac{1}{2} [v(I)^2 + \phi^2] Y^2 P_{YY}(Y, I) + \frac{1}{2} \sigma^2 I^2 (1 - I)^2 P_{II}(Y, I) \\
& + P_{IY} \sigma I (1 - I) v(I) Y + \lambda^\mathbb{Q} [P(e^{n(I)} Y, 0) - P(Y, I)]. \quad (\text{C.63})
\end{aligned}$$

By using the homogeneity property, $P(Y, I) = p(I)Y$, and substituting this expression into (C.63), we obtain the ODE (49) for $p(I)$. Using the first-order condition (FOC) for the HJB equation (C.63), we obtain equation (28) for the optimal mitigation $x(I)$.

References

- Allen, L.J. and Burgin, A.M., 2000. Comparison of deterministic and stochastic SIS and SIR models in discrete time. *Mathematical biosciences*, 163(1), pp.1-33.
- Alvarez, F.E., Argente, D. and Lippi, F., 2020. A simple planning problem for covid-19 lockdown (No. w26981). National Bureau of Economic Research.
- Andersen, A.L., Hansen, E.T., Johannesen, N., Sheridan, A., 2020. Pandemic, shutdown and consumer spending: Lessons from Scandinavian responses to COVID-19, Working Paper.
- Andersson, H., and Britton, T. 2012. *Stochastic epidemic models and their statistical analysis* (Vol. 151). Springer Science and Business Media.
- Atkeson, A., 2020. What will be the economic impact of COVID-19 in the US? Rough estimates of disease scenarios (No. w26867). National Bureau of Economic Research.
- Black, F. and Scholes, M., 1973. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3), pp.637-654.
- Brauer, F., Driessche, P.D. and Wu, J., 2008. *Lecture Notes in Mathematical Epidemiology*. Berlin, Germany: Springer.
- Clancy, D., 2005. A stochastic SIS infection model incorporating indirect transmission. *Journal of applied probability*, 42(3), pp.726-737.
- Cochrane, J.H., 2009. *Asset Pricing: Revised Edition*. Princeton university press.
- Duffie, D., 2001. *Dynamic Asset Pricing Theory*. Princeton University Press.
- Dureau, J., Kalogeropoulos, K. and Baguelin, M., 2013. Capturing the time-varying drivers of an epidemic using stochastic dynamical systems. *Biostatistics*, 14(3), pp.541-555.
- Eichenbaum, M.S., Rebelo, S. and Trabandt, M., 2020. The macroeconomics of epidemics (No. w26882). National Bureau of Economic Research.
- Elenev, Vadim, Landvoigt, Tim, and Van Nieuwerburgh, Stijn, 2020. Can the COVID bailouts save the economy? (no. w27207), National Bureau of Economic Research.
- Farboodi, Maryam, Jarosch, G., and Shimer, R., 2020. Internal and external effects of social distancing in a pandemic. Working Paper.
- Federal Reserve Board, *Financial Stability Report*, May 15, 2020.
- Ferguson, N., Laydon, D., Nedjati Gilani, G., Imai, N., Ainslie, K., Baguelin, M., Bhatia, S., Boonyasiri, A., Cucunuba Perez, Z.U.L.M.A., Cuomo-Dannenburg, G. and Dighe, A., 2020. Report 9: Impact of non-pharmaceutical interventions (NPIs) to reduce COVID-19 mortality and healthcare demand.
- Gormsen, N.J. and Koijen, R.S., 2020. Coronavirus: Impact on stock prices and growth expectations. University of Chicago, Becker Friedman Institute for Economics Working Paper, (2020-22).

- Gourinchas, P.O., 2020. Flattening the pandemic and recession curves. Mitigating the COVID Economic Crisis: Act Fast and Do Whatever, Working Paper, p.31.
- Gray, A., Greenhalgh, D., Hu, L., Mao, X. and Pan, J., 2011. A stochastic differential equation SIS epidemic model. *SIAM Journal on Applied Mathematics*, 71(3), pp.876-902.
- Hong, H., Kubik, J., Wang, N., Xu, X., and Yang, J., 2020. Pandemics, vaccines and corporate earnings. SSRN Working Paper.
- Hong, Harrison, Wang, N. and Yang, J., 2020. Mitigating disaster risks to sustain growth (no. w27066). National Bureau of Economic Research.
- Imai, N., Cori, A., Dorigatti, I., Baguelin, M., Donnelly, C.A., Riley, S. and Ferguson, N.M., 2020. Report 3: transmissibility of 2019-nCoV. Reference Source.
- Kermack, W.O. and McKendrick, A.G., 1927. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 115(772), pp.700-721.
- Kruse, T. and Strack, P., 2020. Optimal control of an epidemic through social distancing. *Cowles Foundation Discussion Paper No. 2229*.
- Kucharski, A.J., Russell, T.W., Diamond, C., Liu, Y., Edmunds, J., Funk, S., Eggo, R.M., Sun, F., Jit, M., Munday, J.D. and Davies, N., 2020. Early dynamics of transmission and control of COVID-19: a mathematical modelling study. *The Lancet Infectious Diseases*.
- Landier, A. and Thesmar, D., 2020. Earnings expectations in the COVID-19 crisis (no. 27160). National Bureau of Economic Research.
- Lucas Jr, R.E., 1978. Asset prices in an exchange economy. *Econometrica: Journal of the Econometric Society*, pp.1429-1445.
- Li, R., Pei, S., Chen, B., Song, Y., Zhang, T., Yang, W. and Shaman, J., 2020. Substantial undocumented infection facilitates the rapid dissemination of novel coronavirus (SARS-CoV-2). *Science*, 368(6490), pp.489-493.
- Manski, C.F. and Molinari, F., 2020. Estimating the COVID-19 infection rate: Anatomy of an inference problem. *Journal of Econometrics*.
- Merton, R.C., 1971. Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory*, 3(4), pp.373-413.
- Sharpe, W.F., 1964. Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance*, 19(3), pp.425-442.
- Zhao, Y. and Jiang, D., 2014. The threshold of a stochastic SIS epidemic model with vaccination. *Applied Mathematics and Computation*, 243, pp.718-727.

Table 3: This table summarizes the parameter values for our epidemic and asset valuation analyses. One period corresponds to one month.

Parameters	Symbol	Value
A. Epidemic and Vaccine		
transmission rate	β	6.616
recovery rate	γ	2.173
volatility of infected population	σ	1.689
arrival rate of vaccine	λ	1.1/12
B. Earnings Process		
earnings growth volatility	ϕ	20%/ $\sqrt{12}$
correlation coefficient	ρ	1
normal-time earnings growth rate	\hat{g}	5%/12
growth reduction level parameter	ζ_1	3
growth reduction curvature parameter	ζ_2	0.25
C. Asset Pricing Parameters		
risk-free rate	r	4%/12
market price of business-as-usual risk	η^B	0.3/ $\sqrt{12}$
market price of pandemic risk	η^Z	-0.4/ $\sqrt{12}$
vaccine arrival timing risk	κ	-1
D. Mitigation Technology		
flow fixed cost	π	0.05
coefficient of the $n(I)$ function	α_1	0.18
curvature of the $n(I)$ function	α_2	0.026
customer impact on transmission	ψ	0.26
firm mitigation efficacy	θ_1	20/12
curvature of firm mitigation efficacy	θ_2	0.032

Table 4: Effects of Vaccine Arrival Rate, λ , on Conditional Distributions of $p(I_t)$ without any Mitigation. In Panels A, B, C, and D, the expected vaccine arrival time $1/\lambda$ is ∞ , six months, one year, and two years. $\alpha_1 = 0$ and other parameter values are given in Table 3.

t	A. No Vaccine ($\lambda = 0$)			B. Six Months ($\lambda = \frac{2}{12}$)		
	Deterministic	Stochastic		Deterministic	Stochastic	
	$p(I_t)$	$\mathbb{E}(p(I_t))$	$\sqrt{\mathbb{V}ar(p(I_t))}$	$p(I_t)$	$\mathbb{E}(p(I_t))$	$\sqrt{\mathbb{V}ar(p(I_t))}$
1 wk	5.527	5.638	0.015	17.565	17.744	0.448
2 wk	5.513	5.626	0.021	17.623	17.799	0.623
4 wk	5.487	5.601	0.029	17.735	17.906	0.855
6 wk	5.461	5.577	0.034	17.844	18.010	1.016
8 wk	5.438	5.553	0.039	17.952	18.109	1.137
3 mo	5.393	5.500	0.041	18.230	18.351	1.331
4 mo	5.384	5.467	0.034	18.486	18.553	1.409
6 mo	5.384	5.439	0.014	18.916	18.933	1.399
9 mo	5.384	5.434	0.002	19.343	19.350	1.213
12 mo	5.384	5.434	0.002	19.601	19.606	0.997
24 mo	5.384	5.434	0.002	19.946	19.947	0.391
∞	5.384	5.434	0.002	20	20	0

t	C. One Year ($\lambda = \frac{1}{12}$)			D. Two Years ($\lambda = \frac{0.5}{12}$)		
	Deterministic	Stochastic		Deterministic	Stochastic	
	$p(I_t)$	$\mathbb{E}(p(I_t))$	$\sqrt{\mathbb{V}ar(p(I_t))}$	$p(I_t)$	$\mathbb{E}(p(I_t))$	$\sqrt{\mathbb{V}ar(p(I_t))}$
1 wk	15.566	15.762	0.591	13.052	13.245	0.664
2 wk	15.616	15.812	0.830	13.088	13.282	0.936
4 wk	15.717	15.911	1.157	13.162	13.356	1.315
6 wk	15.818	16.010	1.396	13.236	13.431	1.600
8 wk	15.920	16.108	1.588	13.314	13.507	1.835
3 mo	16.200	16.357	1.944	13.538	13.706	2.298
4 mo	16.485	16.583	2.149	13.783	13.897	2.598
6 mo	17.026	17.062	2.366	14.281	14.335	3.019
9 mo	17.684	17.706	2.425	14.953	14.993	3.377
12 mo	18.196	18.213	2.342	15.546	15.581	3.559
24 mo	19.336	19.343	1.661	17.299	17.320	3.513
∞	20	20	0	20	20	0