Vaccine Progress, Stock Prices, and the Value of Ending the Pandemic *

Viral V. Acharya† Timothy Johnson‡ Suresh Sundaresan§ Steven Zheng¶

February 2023

Abstract

One measure of the ex ante cost of disasters is the welfare gain from shortening their expected duration. We introduce a stochastic clock into a standard disaster model that summarizes information about progress (positive or negative) towards disaster resolution. We show that the stock market response to duration news is essentially a sufficient statistic to identify the welfare gain to interventions that alter the state. Using information on clinical trial progress during 2020, we build contemporaneous forecasts of the time to vaccine deployment, which provide a measure of the anticipated length of the COVID-19 pandemic. The model can thus be calibrated from market reactions to vaccine news, which we estimate. The estimates imply that ending the pandemic would have been worth from 5% to 15% of total wealth as the expected duration varied in this period.

JEL Codes: D6, D8, E21, E32, Q54

Keywords : COVID-19, welfare costs, stock prices, vaccine news

*First draft: November 15, 2020. We thank Dick Berner, Rob Engle, Lars Lochstoer, Stavros Panageas, Matt Richardson, Venky Venkateswaran, and Olivier Wang for their comments and suggestions. We also thank seminar participants at NYU Stern, NYU Stern Quantitative Finance and Econometrics, Advisory Board Meeting of NYU Stern Volatility and Risk Institute, UIUC Gies, IMF, Shanghai Advanced Institute of Finance, UCLA Anderson, and JHU Carey Finance Conference for valuable comments and suggestions. We are grateful to the Vaccine Centre at the London School of Hygiene & Tropical Medicine for sharing publicly available data from their vaccine development tracker. Our vaccine progress indicator (VPI) is available at giesbusiness.illinois.edu/tcj/data/VPI.xlsx. An earlier version of this paper circulated under the title "The Value of a Cure: An Asset-Pricing Perspective."

†Corresponding author. New York University, Stern School of Business, NBER, and CEPR: vacharya@stern.nyu.edu.

‡University of Illinois at Urbana-Champaign: tcj@illinois.edu.

§Columbia University, Graduate School of Business: ms122@columbia.edu.

¶University of California, Berkeley: steven_zheng@berkeley.edu.

1 Introduction

Quantifying the economic consequences of disasters is a crucial step in assessing policy responses along social, medical, fiscal, and monetary dimensions. One measure of the *ex ante* cost of a disaster is the welfare gain from shortening its expected duration. This paper presents estimates of that quantity for the COVID-19 pandemic during 2020. In the context of a standard disaster model with uncertain duration, we show that the stock market response to duration news effectively identifies the welfare implications of state transitions. We combine this observation with novel data on the progress of vaccine development during 2020, which we use to construct a real-time forecast series for the expected time until successful vaccine deployment. Based on the stock market sensitivity to changes in these expectations, we estimate that ending the pandemic would have been worth 5-15% of total wealth over the sample period.

Global health planners have recently issued several detailed proposals for preparing for future pandemics. (See Craven et al. (2021), Summers et al. (2021), Lander and Sullivan (2021), Gates (2022).) Prominent in each of these has been recommendations for shortening the anticipated time for future vaccine development and deployment, e.g. to within six months of an infectious outbreak. Our paper contributes a novel measurement strategy that can speak to the economic benefits of proposed infrastructure investments to achieve such targets. Our exercise can be viewed as complementary to the standard approach in health economics of computing the cost/benefits of interventions such as vaccines via combining their forecasted effects on infections and deaths with some estimates of (or assumptions about) the monetary value of these outcomes. (We compare our results to some of these estimates in Section 4.)

The welfare calculation we undertake is directly analogous to the seminal work of

Lucas (1987) in assessing the costs associated with business cycle risk. Just as that paper provides a framework for assessing the consequences of policy responses to mitigate consumption volatility, our work speaks to the cost-benefit analysis in alleviating the threat of current and future pandemics. The paper thus contributes to the literature that assesses the welfare costs of disaster risk (see Barro (2009), Martin (2008), Pindyck and Wang (2013), Martin and Pindyck (2015), Jordà et al. (2020), Martin and Pindyck (2021), and Hong et al. (2022)). Uniquely to this literature, we exploit novel information within the crisis itself to identify key quantities in the model: a measure of the expected duration of the pandemic and estimates of the stock market sensitivity to changes in that expectation.

The paper's analysis proceeds as follows.

We start with a standard regime-switching model of repeated disasters. In keeping with a common modeling assumption in the rare disaster literature (see (Barro, 2006; Gabaix, 2012; Gourio, 2012; Tsai and Wachter, 2015)), the nature of a disaster is that agents are exposed to a shock that destroys part of the economy's stock of wealth. Within the disaster, we model sub-regimes corresponding to observable progress towards resolution of the episode. The current state thus provides a stochastic clock within the crisis. The value of a claim to the economy's output responds to the clock transitions because progress out of the disaster regime lowers both the expected wealth losses and uncertainty. Crucially, the magnitude of this response effectively identifies the set of parameters necessary for the welfare calculation. Hence, conditional on this sensitivity, our results are not driven by specific assumptions about the course of the pandemic or about agents' preferences.

Next, we exploit the fact that, during 2020, it was widely believed that deployment of an efficacious vaccine for COVID-19 was both a necessary and sufficient condition for a robust economic recovery.[1] During this period, there was little progress on alternative

[1]As Federal Reserve Chairperson Jay Powell remarked to CBS News in May 2020: "For the economy to fully recover, people will have to be fully confident. And that may have to await the arrival of a vaccine".

mechanisms to resolve the pandemic (e.g. treatments or natural, herd immunity). Moreover, views at the time were optimistic about distribution and uptake of an efficacious vaccine, and the emergence of variants was not yet widely anticipated. Hence, during our sample period, forecasts about vaccine deployment could be viewed as forecasts for the duration of the crisis.[2]

We therefore construct a time-series forecasts for the expected time to the widespread deployment of a successful vaccine that we call the vaccine progress indicator (VPI). The forecasts are based on the chronology of stage-by-stage progress of individual vaccine candidates (obtained from the Vaccine Centre at the London School of Hygiene & Tropical Medicine) and related news (obtained from FactSet). Each day, we assign probabilities to each active candidate's future transition across developmental phases, or to failure. We then simulate these future transitions for all of the candidates, and each run of the simulation produces a first-to-succeed candidate.[3] Averaging across runs of the simulation gives us that day's forecast for the expected time remaining until vaccine deployment. The evolution of our forecast during 2020 is shown in Figure 1. We compare its levels to those of some published contemporaneous survey forecasts in Section 3.2.

We then relate stock market valuations to the expected time to vaccine deployment by regressing returns on changes in our forecast series. Controlling for large market moves attributable to release of other macroeconomic and policy-related news, we estimate that a reduction in the expected time to deployment of a vaccine by one year results in an aggregate stock market return of approximately 5%.[4] The joint relationship exhibits the right cross-sectional properties, with the co-movement between returns and changes in the vac-

[2]Note that our maintained hypothesis that investors believed that vaccine deployment would effectively end the pandemic biases <u>downward</u> our estimate of the value of ending the pandemic because a more skeptical market would react less strongly to vaccine news. So, holding the observed response fixed, the less success a vaccine is expected to have, the greater must be the value of complete success.

[3]An analogy from credit risk literature is that of a first-to-default basket in which several correlated firms are part of a basket and the quantity of interest is the expected time to a first default.

[4]We can and do assess the hypothesis that the market overreacts to vaccine news and that wealth gains subsequently reverse. At short horizons, at least, the evidence suggests the opposite.

Figure 1: Expected Time to Vaccine Deployment

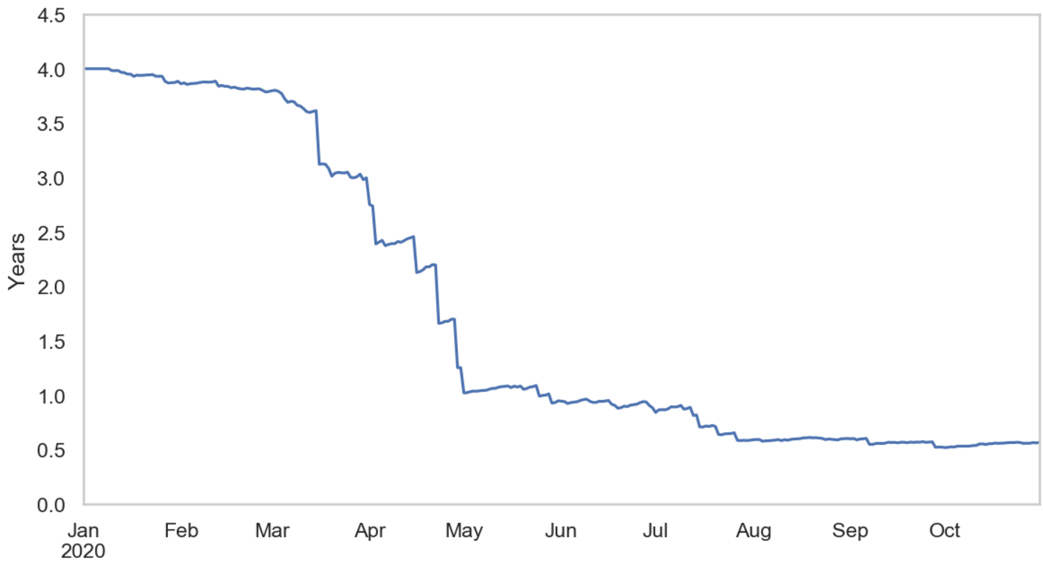

Note: Figure shows our estimate of the expected time to widespread deployment of a COVID-19 vaccine in years.

cine expectation series being stronger for sectors most affected by the onset of COVID-19. This lends support to the conclusion that our series is indeed measuring changes in beliefs about the duration of the pandemic.

Armed with these estimates, we return to the model. Assuming that investors' model of the pandemic economy aligns with the model's dynamics, and under standard preference assumptions, the observed stock market sensitivity implies a welfare cost per year of expected time remaining of approximately the same magnitude as the stock market response coefficient. That is, in late April 2020 when our forecasts imply an expected remaining duration of about one year, the welfare value of ending the pandemic was approximately 5% of total wealth. At the beginning of March, the analogous number is approximately 15%. We discuss interpretation of these magnitudes in Section 4.2.1.

We extend the model in several directions in the online appendix. A first extension endogenizes the real option to invest in vaccine research so that the speed of progress is an equilibrium outcome. A second generalization endogenizes the magnitude of health shocks via labor choice in order to study the role of exposure externalities. Last, the appendix discusses to what extent the model can offer a plausible depiction of the path of the stock market and consumption that was actually experienced during 2020. We argue that, although there are shortcomings, the primary contours of the data can be reasonably described as an outcome within the model, given certain realizations of the stochastic shocks.

To summarize, the contribution of the paper is to bring novel data observed during 2020 to bear on the important question of the *ex-ante* cost of such pandemics. We show that two key quantities that we estimate – the expected duration of the crisis and the stock market response to vaccine progress – are effectively sufficient to identify the welfare gain to interventions curtailing the pandemic.

As in any such exercise, there are inherent limitations in interpreting the results. The model provides only a reduced form depiction of the pandemic, of the economy, and of the stock market. We acknowledge that, in reality, the value of ending the pandemic depended on factors that we omit, including the state of the outbreak (as in SIR-type models), the infectiousness of the virus and its variants, and the stance of fiscal and monetary policy. Without adding more state variables, our theory cannot provide a complete account of global health, economic conditions, or stock market behavior during 2020. It is worth remarking, however, that our methodology and conclusions do not need to assume that the observed stock market responses were perfectly rational, nor that investors *correctly* anticipated either vaccine deployment or the subsequent economic impact.

1.1 Related Literature

As noted above, our approach to quantifying the cost of the pandemic parallels that of the literature assessing the cost of business cycles. While Lucas (1987) finds small welfare improvements to reducing consumption uncertainty, Tallarini Jr (2000) shows that this conclusion is overturned in models with recursive utility when calibrated to match asset pricing moments. Echoing this finding and foreshadowing our own, Barro (2009) reports that, in a model with rare disasters, moderate risk aversion, and an elasticity of intertemporal substitution greater than one, society would willingly pay up to 20% of permanent income to eliminate disaster risk. A number of papers, including Pindyck and Wang (2013), explore the welfare costs associated with climate risk. The latter work addresses the issue of how much should society be willing to pay to reduce the probability or impact of a catastrophe.

A different approach to determining agent's valuation of alternative consumption paths is presented by Alvarez and Jermann (2004) who define the "marginal cost" of business cycles as the ratio of a market price of a claim to the true consumption process to that of an alternative path with the same mean but lower uncertainty. While this approach has the advantage of being preference-free, it is not obvious how to attain the required market prices. (It may well be applicable in the future if pandemic insurance becomes widely traded.)

Our model is related to Ai (2010) and Gourio (2012). We share many of features of each, but differ by offering a setting that allows us to connect to our unique empirical data. Both of those papers feature shocks to the stock of productive capital, with endogenous consumption. Like Ai (2010), our model is in continuous time and allows for frictionless adjustment to capital. That paper also studies the effect of parameter uncertainty, although not in a setting with rare disasters. We differ from Gourio (2012) in modelling *extended* disasters, like pandemics, whose *expected duration* is a crucial and evolving feature of the economy.

While the literature studying the economic impact of COVID-19 has exploded in a short period of time, there is relatively little focus on the role played by vaccine development and its progress.

Hong et al. (2021b) study the effect of pandemics on firm valuation by embedding an asset pricing framework with disease dynamics and a stochastic transmission rate, equipping firms with pandemic mitigation technologies. Similar to our paper, they model vaccine arrival as a Poisson jump process between pandemic and non-pandemic states. Hong et al. (2021a) combine the model of Hong et al. (2021b) with pre- and post-COVID-19 analyst forecasts to infer market expectations regarding the arrival rate of an effective vaccine and to estimate the direct effect of infections on growth rates of earnings. In particular, they develop a regime-switching model of sector-level earnings with shifts in their first and second moments across regimes.

In both of these papers, the pricing kernel is exogenously specified for the pandemic and the non-pandemic states. In contrast, our model is general equilibrium in nature with the representative agent optimally choosing labor and consumption (and, in turn, investment in capital) to mitigate health risk. Deriving asset prices from first principles in a regime-switching framework of with news about transition probabilities is a theoretical contribution of our paper.

Elenev et al. (2020) incorporate a pandemic state with low, disperse firm productivity that recurs with low probability for studying government intervention in corporate credit markets. Hong et al. (2021b) fix expected pandemic duration around one year but show in comparative statics that asset prices can have considerable sensitivity to the arrival rate of the vaccine. Hong et al. (2021a) use their model to infer the arrival rate of the vaccine. In contrast, we construct an estimated time to vaccine deployment using news on the progress of clinical trials of vaccines for COVID-19. We infer the loss in economic wealth in the pandemic from stock market reactions to changes in these forecasts

On the empirical side, Baker et al. (2020a) deploy transaction-level data to study consumption responses to COVID-19, finding an increase in the beginning in an attempt to stockpile home goods, followed by a sharp decrease as the virus spread and stay at home orders were enforced. Using customized survey data, Coibion et al. (2020) find lockdowns decreased consumer spending by 30 percent, with the largest drops in travel and clothing. Bachas et al. (2020) find a rebound in spending, especially for low-income households, since mid-April. Chetty et al. (2020) further find high-income households significantly reduced spending, especially on services that require in-person interactions, leading to business losses and layoffs in the most affluent neighborhoods. Outside the US, Sheridan et al. (2020) and Andersen et al. (2020) find aggregate spending decreased 27% in the first seven weeks following Denmark's shutdown, with the majority of the decline caused by the virus itself regardless of social distancing laws. Chen et al. (2020) use daily transaction data in China and find severe declines in spending, especially in dining, entertainment and travel sectors. In Section 4.2.1 we relate our estimates and methodology to the broader literature attempting to measure the costs associated with COVID-19. Like us, Del Angel et al. (2021) measure stock market sensitivity to pandemic-related news. They find a strong negative correlation between reported fatalities and market returns during the 1918 flu outbreak.

2 Disasters with News about Duration

To begin, we modify a standard rare-disaster framework to incorporate news about the expected duration of disaster episodes. We show how to compute welfare in terms of the economy's state. Then we price a claim to output, which is a function of the expected duration. The key finding is that that functional relationship effectively pins down the magnitude of welfare effects. This lowers the effective dimensionality of the model, enabling a feasible calibration whose conclusions are largely insensitive to the particular parameters governing the disaster dynamics and household preferences.

2.1 Disaster Dynamics

Our underlying building block is a general equilibrium regime switching model of disasters (following Nakamura et al. (2013) and Collin-Dufresne et al. (2016)) where the economy is alternately in "normal times" and in a "disaster" regime. The regimes differ in their state-specific stochastic process for the accumulation of wealth. Fundamentally, the model depicts disasters as destroying or degrading that stock of wealth, as in Gourio

(2012), with consumption and other policies potentially responding endogenously. For this reason, we work with a production-based framework rather than an endowment economy.

Specifically, let q denote the quantity of productive capital of an individual household (which could be viewed as both physical and human capital). We assume that the stock of q is freely convertible into a flow of consumption goods at rate C per unit time. Then our specification is that, in state s, q evolves according to the process

$$dq = \mu(s)qdt - C(s)dt + \sigma(s)qdB_t - \chi(s)qdJ_t \tag{1}$$

where B_t is a standard Brownian Motion and J_t is a Poisson process with intensity $\zeta(s)$. We set $\chi(s) = 0$ in normal times and $\chi(s) > 0$ within the disaster. Hence we interpret the Poisson shock as capturing the risk of an economic loss due to being in a disaster regime.

Next, we augment the disaster regime to include sub-regimes that correspond to observable progress to returning to normal life. Now let the state s take values in $\{0, 1, \ldots, S-1, S\}$, where both $s = 0$ and $s = S$ are the same normal time regime, and the others are the disaster sub-states. We assume that the economy switches between these states based on a Markov-switching or transition matrix. The transition probabilities are as follows:

$$Pr(s_{t+dt} = 1 | s_t = 0 \text{ or } S) = \eta dt \tag{2}$$

$$Pr(s_{t+dt} = s_t | s_t = 0 \text{ or } S) = 1 - \eta dt \tag{3}$$

$$Pr(s_{t+dt} = s-1 | s_t = s \in [1, S-1]) = \lambda_d(s)dt \tag{4}$$

$$Pr(s_{t+dt} = s+1 | s_t = s \in [1, S-1]) = \lambda_u(s)dt \tag{5}$$

$$Pr(s_{t+dt} = s_t | s_t = s \in [1, S-1]) = 1 - \lambda_d(s)dt - \lambda_u(s)dt. \tag{6}$$

That is, η is the intensity of disaster arrivals, and λ_d and λ_u are the respective intensities within a disaster for transitions “down" or “up" to the adjacent states. Given this specification, a straightforward Markov chain calculation yields $\mathbb{E}_t[T^\star|s]$ where $T^\star$ is the time at which the state S is attained and the current disaster is terminated.

Although the specification allows for arbitrary parameter differences across the sub-states , our intention is rather to interpret them as differing *only* in so far as advances in the state reduce the expected time to exit the disaster. Hence, for $0 < s < S$, we will take $\mu(s) = \mu(1)$, $\sigma(s) = \sigma(1)$, and $\chi(s) = \chi, \zeta(s) = \zeta$ to all be constants. We can thus interpret the state index s as a probabilistic clock. As such, within a disaster, agents are subject to an additional source of risk: stochastic duration.[5]

It is worth highlighting the implicit assumptions in the model about long-run effects. Our specification is pessimistic in the sense that loss of wealth due to the disaster shocks is permanent. Productive capital q does not get restored when the disaster ends. On the other hand, the model is optimistic in the sense that the productive *process*, dq, does fully revert to pre-disaster dynamics. After the disaster, the world looks stochastically the same as it did before. Agents know that the disaster will eventually end, and there will be no scarring effects, e.g., on the economy's normal growth rate, $\mu(0)$.

2.2 Agents

We assume the economy has a unit mass of identical agents (households). Each agent has stochastic differential utility or Epstein-Zin preferences (Duffie and Epstein, 1992; Duffie and Skiadas, 1994) based on consumption flow rate C, given as

$$\mathbb{J}_t = \mathbb{E}_t \left[\int_t^\infty f(C_{t'}, \mathbb{J}_{t'}) dt' \right] \tag{7}$$

[5] In a business cycle context, Andrei et al. (2019) model an economy in which the mean-reversion speed of current consumption shocks is time-varying, and thus agents similarly face persistence risk.

and aggregator

$$f(C,\mathbb{J}) = \frac{\rho}{1-\psi^{-1}}\left[\frac{C^{1-\psi^{-1}} - [(1-\gamma)\mathbb{J}]^{\frac{1}{\theta}}}{[(1-\gamma)\mathbb{J}]^{\frac{1}{\theta}-1}}\right] \tag{8}$$

where $0 < \rho < 1$ is the discount factor, $\gamma \geq 0$ is the coefficient of relative risk aversion (RRA), $\psi \geq 0$ is the elasticity of intertemporal substitution (EIS), and

$$\theta^{-1} \equiv \frac{1-\psi^{-1}}{1-\gamma} \tag{9}$$

The use of recursive preferences is standard in macrofinance models because of their ability to match financial moments. We recognize the limitations of using a utility specification driven by consumption goods, particularly within a crisis when other considerations (e.g., health, social interaction, the safety of others) so strongly affect well-being. However, using a familiar formulation ensures that our findings are not driven by non-standard assumptions about utility. The representative agent's problem is, in each state s, to choose optimal consumption $C(s)$ that maximizes the objective function $\mathbb{J}(s)$.

2.3 Solution

We now characterize the solution to the optimization problem.

Proposition 1. *Denote*

$$g(s) \equiv \frac{(1-\gamma)\rho}{(1-\psi^{-1})} - (1-\gamma)\left(\mu(s) - \frac{1}{2}\gamma\sigma(s)^2\right) - \left([1-\chi(s)]^{1-\gamma} - 1\right) \tag{10}$$

Let $H(s)$'s denote the solution to the following system of S recursive equations:

$$g_0 \equiv g(0) = \frac{(1-\gamma)}{(\psi-1)}\rho^{\psi}\left(H(0)\right)^{-\psi\theta^{-1}} + \eta\left[\frac{H(1)}{H(0)} - 1\right] \tag{11}$$

$$g_1 \equiv g(1) = \frac{(1-\gamma)}{(\psi-1)}\rho^{\psi}\left(H(s)\right)^{-\psi\theta^{-1}} + \lambda_d\left[\frac{H(s-1)}{H(s)} - 1\right] + \lambda_u\left[\frac{H(s+1)}{H(s)} - 1\right], \tag{12}$$

for $s \in \{1, \ldots, S-1\}$.

Assuming the solutions are positive, optimal consumption in state s is

$$C(s) = \frac{\left(H(s)\right)^{-\psi\theta^{-1}} q}{\rho^{-\psi}}, \tag{13}$$

and the value function of the representative agent is

$$\mathbb{J}(s) \equiv \frac{H(s)q^{1-\gamma}}{1-\gamma} \tag{14}$$

Note: All proofs appear in the appendix.

The recursive system is straightforward to solve numerically. Henceforth we implicitly assume the parameters are such that a unique solution vector $H(s)$ exists and is strictly positive.

2.4 Welfare Across States

The certainty equivalent change in the representative agent's lifetime value function upon a transition from state s to state 0 (or to state S) is given by:

$$V(s) \equiv 1 - \left(\frac{H(s)}{H(0)}\right)^{\frac{1}{1-\gamma}}. \tag{15}$$

This is the percentage of the agent's stock of wealth q that, if surrendered, would be fully compensated by the utility gain of reverting to the non-disaster state. This willingness-to-pay definition is standard in the literature. Using the optimal consumption characterized above, we also obtain that

Proposition 2. *The value of ending the disaster in state s is determined by the ratio of marginal propensity to consume* ($c \equiv dC/dq$) *in the disaster state s relative to that in the non-disaster state, adjusted by the agent's elasticity of intertemporal substitution (EIS):*

$$V(s) = 1 - \left(\frac{c(s)}{c(0)}\right)^{-\frac{1}{\psi-1}} = 1 - \left(\frac{C(s)}{C(0)}\right)^{-\frac{1}{\psi-1}} \tag{16}$$

The definition above naturally generalizes to the willingness-to-pay to alter the clock from any given s to another s', as

$$V(s,s') \equiv 1 - \left(\frac{H(s)}{H(s')}\right)^{\frac{1}{1-\gamma}}. \tag{17}$$

This is the relevant quantity when assessing mitigations that, while not ending the disaster, shorten its remaining expected duration.

2.5 Asset Pricing

We interpret "the market portfolio" within the model as a claim to the economy's output.[6] Output is the net new resources per unit time, which is implicitly defined by two endogenous quantities: the change in the cumulative wealth plus consumption, or $dq + Cdt$. Denote the price of the output claim as $P = P(s,q)$. By the fundamental theorem of asset pricing, the instantaneous expected excess return to holding this claim is equal to minus the covariance of its returns with the pricing kernel. Under stochastic differential utility,

[6] Note that this is not the same as a claim to aggregate consumption. As is well known (see Ai (2010)), in an economy where the capital stock can be costlessly converted to consumption goods, the consumption claim's price is equal to the capital stock, q. However, any wedge between consumption and payouts to equity will result in a nonconstant price-capital ratio. Our assumption here is that the expected cash flow to the market portfolio mirrors the expected impact of the disaster on wealth. In the online appendix, we describe a decentralization of the economy in which this cash flow is the net payout of the corporate sector to households.

and with the value function solution above, the pricing kernel in our economy is given by

$$\Lambda_t = \exp\left\{\int_0^t \left[\rho^{\psi}(\theta/\psi)H(s_u)^{-\psi/\theta}\right] du\right\} q_t^{-\gamma} H(s_t)$$

From this, we derive the value of the market portfolio in the following proposition.

Proposition 3. *The price of the output claim is $P = p(s)q$ where the constants $p(s)$ solve a matrix system $Y = Xp$ where X is an $S+1$-by-$S+1$ matrix and Y is an $S+1$ vector both of whose elements are given in the appendix.*

Henceforth we assume the model parameters are such that the matrix X defined in the proposition is of full rank. The behavior of the price-capital ratio, $p(s)$, accords with economic intuition: it declines sharply on a move from state $s = 0$ to $s = 1$, and then gradually (and approximately linearly) recovers as s advances. Thus, the quantity $\Delta \log P = \log p(s+1) - \log p(s)$ is positive for $s > 0$ and, in practice, varies little with s.

2.6 Welfare and Stock Market Sensitivity to Duration News

Next, define $T^\star$ as the time at which the state S is attained and the disaster is terminated. In the next section, we construct an empirical counterpart to its time t expectation, $\mathbb{E}_t[T^\star]$, during 2020. It is straightforward to show that this expectation is again given by a linear system, which we omit for brevity. Moreover, for large S, the difference

$$\Delta\mathbb{E}[T^\star] = \mathbb{E}[T^\star|s+1] - \mathbb{E}[T^\star|s] \sim \frac{1}{\lambda_u} \tag{18}$$

is effectively constant as well. Given $P(s)$ and $\mathbb{E}[T^\star]$, we can readily compute the sensitivity

$$\frac{\Delta \log P}{\Delta\mathbb{E}[T^\star]}. \tag{19}$$

This corresponds to the second quantity we attempt to measure empirically.

The goal of the paper is to estimate the quantities (15) and (17), which clearly depend on how the model is calibrated. An important insight, however, is that the model counterpart to the stock market sensitivity restricts the parameters governing the severity of the disaster: $\mu(1), \sigma(1), \chi$, and ζ. More precisely, sets of parameters consistent with an observed value of (19) all have welfare costs, V, within a tight range. This is illustrated in the left-hand panel of Figure 2. The figure uses the baseline parameters shown in Table 1 and varies the disaster parameters over broad ranges.[7] The figure plots the welfare cost per unit time, $V/\mathbb{E}[T^{\star}]$, as a function of the quantity (19), with each point corresponding to a parameter configuration. The key finding is that the vertical spread across model solutions for a given value on the horizontal axis is narrow. This implies that our welfare conclusions, conditional on (19), do not depend upon the precise assumptions about these parameters.[8] While this is a numerical result, the economic logic underpinning it is solid. Although the household's value function and the stock price are different economic constructs, the effect of a disaster on each is largely driven by the same two things: the decrease in expected rate of wealth accumulation (or, the severity of the disaster), and the expected time until a return to normal conditions (the expected duration of the disaster).

Moreover, the same logic implies the welfare results are also not strongly sensitive to the choice of the baseline parameters in Table 1. The parameters in the table are chosen to be consistent with the macro-finance literature (and the rare disaster literature in particular). And, while unconditionally, the welfare cost of a disaster will be strongly affected by, e.g., the assumed degree of risk aversion, γ, the middle panel of Figure 2 shows that raising γ from 4 to 6 largely preserves the identification result from the baseline case. The curve has shifted modestly downward, meaning the disaster is actually slightly less costly in welfare terms once the disaster parameters have been adjusted to

[7] The parameter ranges are given in the figure caption.

[8] The solutions take the number of states to be $S = 12$, which is arbitrary. Results are not sensitive to the choice of the number of states, given the unconditional expected duration of the pandemic. Hereafter we denote λ_u/S as λ without a subscript. We also set the intensity of regress to be $\lambda_d = 0$.

fix the value (19). Similar remarks apply to variation in the other preference parameters and the non-disaster output parameters.[9] The right-most panel varies the baseline disaster timing parameters to correspond to more frequent, shorter disasters. (Specifically, the plot uses $\eta = 0.04, \lambda = 0.5$.) These imply slightly larger welfare costs conditional on (19). But again the curve is narrow: the precise choices of the disaster severity parameters do not affect the welfare conclusions.

Figure 2: Stock Market Sensitivity and Welfare Loss Rate

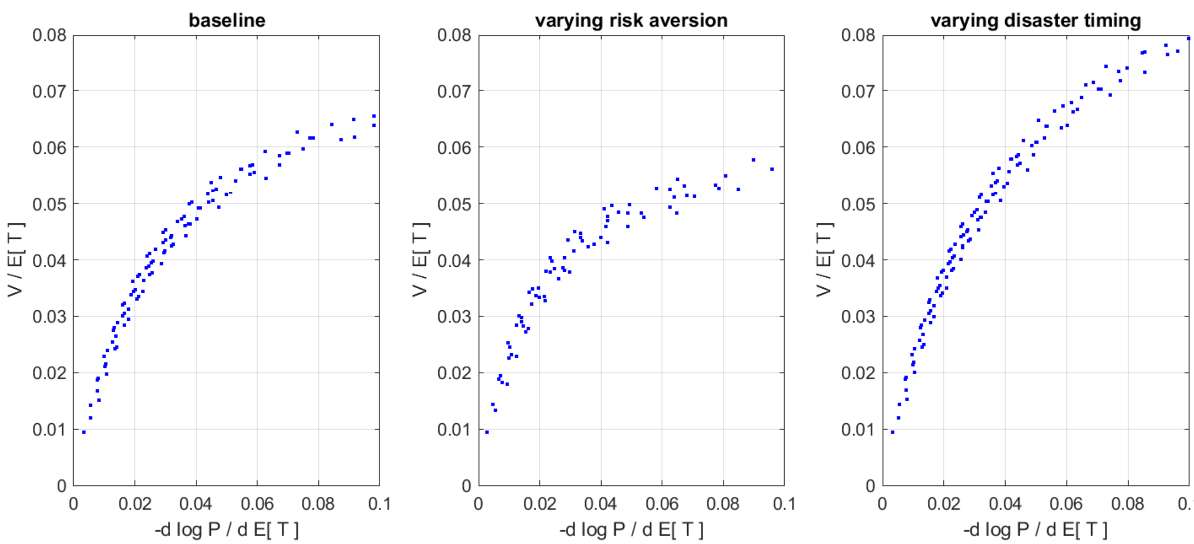

The Figure shows the welfare cost per unit (expected) time, $V/\mathbb{E}[T^{\star}]$ as a function of the stock market sensitivity to changes in the expected time $-\Delta \log P/\Delta\mathbb{E}[T^{\star}]$ as the current state s increases by 1. Each point corresponds to a different set of disaster parameters. The ranges of these parameters are $\mu(1) \in [00.03], \sigma(1) \in [0.050.075], \chi \in [0.0050.05], \zeta \in [12]$. In the left-hand panel the remaining parameters are those given in Table 1. The middle panel uses $\gamma = 6, \mu(0) = 0.035$. The right panel uses $\eta = 0.04, \lambda = 0.5$.

3 Vaccine Progress in 2020

We now turn to the application of our model to the COVID-19 pandemic. As discussed in the introduction, our maintained hypothesis is that expectations about remaining the time to successful deployment of a vaccine during 2020 map into our model's expectation

[9] The results do require an EIS greater than 1. Otherwise the stock market declines on positive duration news. The curves are also sensitive to the normal time riskless rate $r(0)$. The middle panel adjusts $mu(0)$ to fix $r(0) = 0.02$.

for the termination of the pandemic and the return to normal economic conditions. We therefore construct an estimator of that quantity using a statistical model of clinical trial progress.

As described in more detail in section 3.1 below, each day, the model simulates the future progress of each candidate in order to forecast the time until the first candidate successfully deploys. Each run of the simulation takes the current stage of each candidate, then simulates forward the entire clinical timeline through pre-clinical trials, clinical trials, application submission, regulatory approval and vaccine deployment.[10] Each clinical stage is characterized by an expected duration and a probability of failing. We update these probabilities as news arrives about individual candidates: preliminary results or more complete information about earlier trials may be published, released to the press, or leaked. At each simulated stage, each candidate can either fail (and that candidate's run ends) or advance onto the next stage (and the simulation continues). The model then records the time to deployment among candidates successfully reaching deployment, and the average time across a large number of runs is that day's estimate of the expected time until a successful deployment.

Note that vaccine deployment is a final stage with a non-zero probability of failure. That is, an approved vaccine possibly still could not attain widespread deployment, e.g., due to manufacturing and distribution difficulty, emergence of serious safety concerns, mutation of the virus, or adoption hesitancy. In other words, observers at the time were aware that the approval of a vaccine would not necessarily correspond to a "magic bullet" that instantly terminates the spread of COVID-19 and ends the pandemic.

The forecast incorporates the expectation that the success and failure of candidates is positively correlated. This positive dependence arises most obviously because all vac-

[10] This is a simplification. Candidate vaccines will actually undergo multiple overlapping trial sequences with different patient populations, delivery modalities, or medical objectives (endpoints). One sequence could fail while others succeed. Trials can also combine phases I and II or II and III. In our empirical implementation we focus on the most advanced trial of a candidate. This follows Wong et al. (2018).

cines were targeting the same pathogen, and would succeed or fail largely due to its inherent biological strengths and weaknesses. Candidates also shared one of a handful of strategies (or platforms) for stimulating immunological response,[11] relied on common technological components, resources or abilities, and some research teams concurrently developed several candidate vaccines. The model incorporates correlation among candidates by assuming the stochastic durations of each stage are generated by a Gaussian copula with positive correlation matrix.

The next subsection provides details on the data and the each step in the construction of the forecast time-series and discusses some of the underlying assumptions. Further detail on specific assumptions are provided in the online appendix.

3.1 Forecast Construction

Each day, we start with N positively correlated vaccine candidates, with correlation matrix $\mathcal{R}$. Each candidate n is in a state $s \in S$, where

$$S = \{\text{failure}, \text{preclinical}, \text{phase 1}, \text{phase 2}, \text{phase 3}, \text{application}, \text{approval}, \text{deployment}\}$$

and each state has expected duration τ_s and baseline probability of success Π_s^{base}.[12]

Next we augment the state-level, baseline probability of successes with candidate-specific news. Let $\omega_{n,t} \in \Omega$ denote news published at time t about candidate n. For example, Ω could span positive data releases, negative data releases, next state announcements, etc. Then let $\Delta\pi :\to [-1,1]$ be a mapping from news to changes in probabilities. For each candidate, we cumulate the changes in probabilities from all news from the be-

[11] For example, if an RNA-based platform proves to be safe and effective, then all candidates in this family would have a higher likelihood of success. In October 2020, two candidate vaccines had their trials paused due to adverse reactions: both were based on adenovirus platforms.

[12] Our baseline success probabilities employ estimates in Pronker et al. (2013) augmented by our own sample of historical outcomes of infectious disease vaccine trials from pharmaceutical research firm BioMedTracker. Our baseline duration estimates are based on projections from the pharmaceutical and financial press during 2020 as detailed in the online appendix.

ginning of our sample t_0 up to time t,

$$\Delta\pi_{n,t}^{\text{news}} = \sum_{t'=t_0}^{t} \Delta\pi\left(\omega_{n,t'}\right).$$

Finally, we combine it with the baseline probability of success, resulting in a candidate-specific probability of success that potentially varies over time, even within the same state,

$$\pi_{n,s,t}^{\text{total}} = \frac{\exp \Upsilon_{n,s,t}}{1+\exp \Upsilon_{n,s,t}}$$

where $\Upsilon_{n,s,t} = \log \frac{\pi_s^{\text{base}}}{1-\pi_s^{\text{base}}} + 2\Delta\pi_{n,t}^{\text{news}}$.

We simulate stage-by-stage progress of each candidate and generate the expected time to first vaccine deployment, similar to a first-to-default credit model. Specifically, on each day, one run of the simulation repeats steps one and two until candidates have all failed or deployed.[13]

1. We model each state transition as a 2-state Markov chain with exponentially distributed times. Draw an N-dimensional multivariate normal random variable

   $$z_t = [z_{1,t}, \ldots, z_{N,t}]' \sim \mathcal{N}(0, \mathcal{R})$$

   and for each candidate, transform to exponential time with intensity $\lambda_{n,s,t} = \frac{\pi_{n,s,t}}{\tau_s}$

   $$t_{n,s,t} = -\frac{\log \Phi(z_{n,t})}{\lambda_{n,s,t}}.$$

2. Then draw a success or failure Bernoulli random variable with parameter π_s. If failure, then that candidate's run is over. Else if success, then that candidate advances states and the run continues.

[13]The simulation procedure is presented graphically in a flowchart in the appendix.

3. Calculate each candidate's time to vaccine deployment as

$$T_n = \begin{cases} \sum_s t_{n,s,t} & \text{, if candidate deploys} \\ \infty & \text{, if candidate fails} \end{cases}$$

4. Then calculate minimum time to vaccine deployment across candidates, $\min_n T_n$.

That finishes one run of the simulation. We repeat for 50,000 runs and take the cross-run average as T^D, before advancing to the next day. On each day across runs, we calculate the average

$$\mathbb{E}[T^*] = (1-\mu)T_t^D + \mu T^{ND},$$

where some fraction, μ, of simulations will result in all candidates not reaching deployment, so we incorporate T^{ND}, an estimated expected time to deployment by a project outside of our sample.

We obtain the pre-clinical dates and trial history of vaccine candidates from publicly available data collated by the London School of Hygiene & Tropical Medicine (LSHTM). We observe the start dates and durations of each pre-clinical and clinical trial, along with their vaccine strategy. We augment the LSHTM timeline with news pertinent to vaccine progress from FactSet StreetAccount. We classify vaccine related stories into seven positive types and six negative types. The Online Appendix includes more detail on the number of candidates and breakdown of strategies, the news types and corresponding probability adjustments, and our choices of parameters, also presenting evidence that our assumptions are reasonably consistent with the (small) set of observed trial outcomes. We will validate our choices both by examining robustness to reasonable variations and by comparing them to other actual *ex ante* forecasts published during the sample period.

Our indicator of vaccine progress aims to capture expectations about deployment principally in the U.S. since this is likely to be the primary concern of U.S. markets. Because of political considerations, we believe that observers at the time judged it to be very improbable that vaccines being developed in China and Russia would be the first to achieve widespread deployment in the U.S. Our base case construction for this reason omits candidates coded in the LSHTM data as originating in Russia or China, retaining candidates coded as multi-country projects including these two countries. We will also verify that including them in our index does not change our primary results.

In focusing on the scientific advancement of the individual candidates, our measure does not attempt to capture general news about the vaccine development environment and policy. News about the acquisition and deployment of delivery infrastructure by governments (or the failure to do so) could certainly affect estimates of the time to availability. We also do not capture the news content of government financial support programs or pre-purchase agreements. News about regulatory approval standards could have affected forecasts as well. While we could alter our index based on some assessment of the impact of news of this type, we feel we have less basis for making such adjustments than we do for modeling clinical trial progress.

Figure 1 shows the model's estimation of the expected time to widespread deployment from January through October of 2020. The starting value of the index, in January, is determined by our choice of the parameter T^{ND} because, with very few candidates and none in clinical trials, there was a high probability that the first success would come from a candidate not yet active. However this parameter effectively becomes irrelevant by March when there are dozens of projects. The index is almost monotonically declining, since there were no reported trial failures and very few instances of negative news through at least August. The crucial aspects of the index for our purposes are the timing and sizes of the down jumps corresponding to the arrival of good news.

3.2 Validation

We are aware of two datasets that contain actual forecasts of vaccine arrival times, as made in real-time during 2020. As a validation check, we compare our index to these.[14]

The two data sets are surveys, to which individuals supplied their forecasts of the earliest date of vaccine availability. Comparisons between these pooled forecasts and our index require some intermediate steps and assumptions. In both cases, the outcomes being forecast are given as pre-specified date ranges. Thus, on each survey date, we know the percentage of respondents whose point forecast fell in distinct bins. For each survey we estimate the median response, assuming a uniform distribution of responses within the bin containing the median. Under the same assumption, we can also tabulate the percentage of forecasters above and below our index.

The first survey is conducted by Deutsche Bank and sent to 800 "global market participants" asking them when they think the first "working" vaccine will be "available". The survey was conducted four times between May and September. The second survey is conducted by Good Judgement Inc., a consulting firm that solicits the opinion of "elite superforecasters." Their question asks specifically "when will enough doses of FDA-approved COVID-19 vaccine(s) to inoculate 25 million people be distributed in the United States?" (Information about the number of responders is not available.) Responses are tabulated daily, starting from April 24th. For brevity, we examine month-end dates. Table 2 summarizes the comparison.

Our forecasts align well with those of the Deutsche Bank survey, though ours are more optimistic than the median. The optimism is more pronounced when compared to the superforecasters early in the pandemic. Although we are within the interquartile range of forecasts after May, the earlier dates see us in the left-tail of the distribution. A possible

[14] We do not employ these in our empirical work because the forecasters may condition their views on contemporaneous stock market reactions, whereas our measure does not employ any financial market information.

reason is the particular survey question, which specifies an exact quantity of the vaccine being distributed in the United States. Respondents may have more skeptical of feasible deployment than we have assumed. We will examine robustness of our results below to increasing the probability of an approved vaccine failing in the deployment stage.

4 Results

This section presents the paper's main results: measuring the stock market sensitivity to vaccine news and interpreting the estimates through the lens of the model in Section 2.

4.1 Market Reaction to Vaccine News

Our empirical methodology is straightforward: we regress daily market returns in 2020 on changes in the VPI. Since our forecast construction does not utilize any stock market or financial information, its changes are exogenous in the regression context.

A important consideration is controlling for other news, of which there was a great deal during this period, principally due to the extent of the pandemic, policy responses, and the likely economic impact of these, but also including, e.g., the U.S. election cycle. Our approach to controlling for non-vaccine news is to exclude days with large stock market moves that were reliably judged to be due to other sources. Specifically, we employ the classification of Baker et al. (2020b) who analyze causes of daily market moves greater than 2.5% in absolute value. Those authors enlist the opinion of three analysts for each such day and ask them to assign weights to *types* of causes (e.g., corporate news, election results, monetary policy, etc). Under their classification, pandemic-related economic and policy news is assigned one of these categories. Research on vaccines falls under the category "other". We view market returns as very unlikely to have been driven by vaccine news if none of the three analysts assigns more than 25% weight to this category, or if the return was more negative than -2.5%. The latter exclusion is based on the fact that there were no significant vaccine setbacks prior to the end of our data window,[15] and on the assumption that positive vaccine progress cannot have been negative news. We then include dummies for all of the non-vaccine large-news days. There are 28 such days, 17 of which were in March. While the approach is imperfect, it avoids putting (endogenous) financial variables, such as bond yields or credit spreads, on the right hand side of our regressions. And, at a minimum we are limiting the ability of our estimation to misattribute the largest market moves to vaccine progress.

Table 3 shows the resulting regression estimates of market impact. The dependent variable is the return to the value-weighted CRSP index from January 1 through October 31, 2020. The regression specifications include contemporaneous changes in the vaccine progress indicator and also examine lagged effects. Given the sheer volume of news being processed during this period, we do not rule out delayed incorporation of information, which would show up in the lag coefficients. Additionally, including lag terms addresses the possibility that the market overreacts to vaccine news. The specifications also include two lags of the dependent variable to control for short-term liquidity effects. Specifically, the regression is

$$R^e_{m,t} = \alpha + \sum_{h=-k}^{0} \beta_h \Delta\mathrm{VPI}_{t+h} + \gamma_1 R^e_{m,t-1} + \gamma_2 R^e_{m,t-2} + \sum_{j=1}^{28} \delta_j \mathbb{1}_{\mathrm{jump}\ j} + \epsilon_t \tag{20}$$

where $\Delta\mathrm{VPI}_t$ is the change in vaccine progress indicator, and $\mathbb{1}_{\mathrm{jump}\ j}$ is a dummy equal to one on the jth jump date from Baker et al. (2020b). We also present results for specifications without the intercept α, the argument being that including it treats both the realized market rally (or, positive sample average return) and the realized vaccine progress (negative average change in VPI) as having been *expected*, which is a potential misspecifica-

[15] As of the time of this draft, Baker et al. (2020b)'s website had classified days through June 2020. We append September 3 and September 23 as two dates with negative jumps that were attributable to non-vaccine related news.

tion.[16]

The first two columns of the table shows results using our baseline vaccine progress indicator, and without lags. The estimated contemporaneous coefficient from the first column implies a stock market increase of 1.29% on a decrease in expected time to vaccine deployment of one year. While economically meaningful, the response is not statistically significant given the short sample. Also, the estimated intercept in this specification illustrates the point above. A positive value of 0.29 translates into an "expected" annual return of 106% ($1.0029^{250} - 1$) for the stock market, which is not plausible. Column two shows the result of dropping the intercept. The response coefficient rises to 2.08% which is statistically significant at the 10% level

Columns three and four add a single lag of the VPI series. We explored a number of lag specifications, and this is the one preferred by the Bayes Information Criterion. The evidence is consistent with a substantial continued positive market response to vaccine progress as information is processed over a second day. Focusing on the cumulative impact of both days, the sum of the βs is 5.24%, or 6.59% without the intercept, and both are statistically significant at the 1% level. Going forward, we will adopt the more conservative of these two as our primary estimate. The final two columns include a longer lag term, namely the cumulative change in VPI from day $t-5$ to $t-2$. This specification allows us to address the possibility that some of the positive market reaction to vaccine progress in the first two days is an overreaction that subsequently reverses. In fact, we find the reverse: the response continues in the same direction. The cumulative 5-day response now rises to 6.81%, or 7.88% without the intercept.

Our baseline estimate of approximately a 5% response to a one-year change in the VPI is strongly supported by the observed reactions to some especially salient announcements. For example, market rallies of 1.1% and 0.9% followed release of Phase I results

[16] Note that the VPI is, by construction, a conditional expectation, whose *ex ante* expected change is therefore zero. Slightly more accurately, if T^* is the forecast deployment time, the date-t VPI may be expressed as $\mathrm{E}_t[T-t]$ whose expected change over a day is minus one day or approximately -0.004 years.

by Moderna on May 18th and July 14th.[17] These announcements caused drops in the VPI of 0.02 and 0.11 years, respectively, implying a response coefficient of at least 8. In an out-of-sample observation,[18] Pfizer's release of Phase III results on November 9th at 6:45AM caused a pre-opening market surge of 2.8%. Even attributing to the news a reduction in VPI of 0.25 years, the implied response coefficient would be over 10.

Since the construction of the VPI forecasts involves a number of assumptions, the Table 4 presents additional regressions using the baseline specification from column 3 for several variants of the methodology, including the altering the news adjustments, the correlation assumptions, and the probability of successful vaccine deployment. In all of these variations, the estimated market response to vaccine progress is similar in magnitude to those reported in Table 3.

4.1.1 Industry Responses

As a validity check for our findings, we examine the price impact of vaccine progress in the cross-section of industries. We first gauge each industry's exposure to COVID-19 by its cumulative return from February 1, 2020 to March 22, 2020. This period captures the rapid onset of COVID-19 in the US, with a public health emergency being declared on January 31, 2020and a national emergency declared on March 13, 2020. Importantly, this period precedes the Federal Reserve's announcement of the Primary Market Corporate Credit Facility and Secondary Market Corporate Credit Facility on March 23, 2020, helping us pin down industry covariances with COVID-19 itself, separate from covariances with policy responses.

We then estimate industry sensitivity to vaccine progress over the non-overlapping

[17] Returns are computed from after-hours S&P500 futures changes in a 60 minute post-announcement window.

[18] Our clinical trial data ends in October 2020. As of the end of that month, the level of VPI was 0.55.

Figure 3: Industry Sensitivity to Vaccine Progress

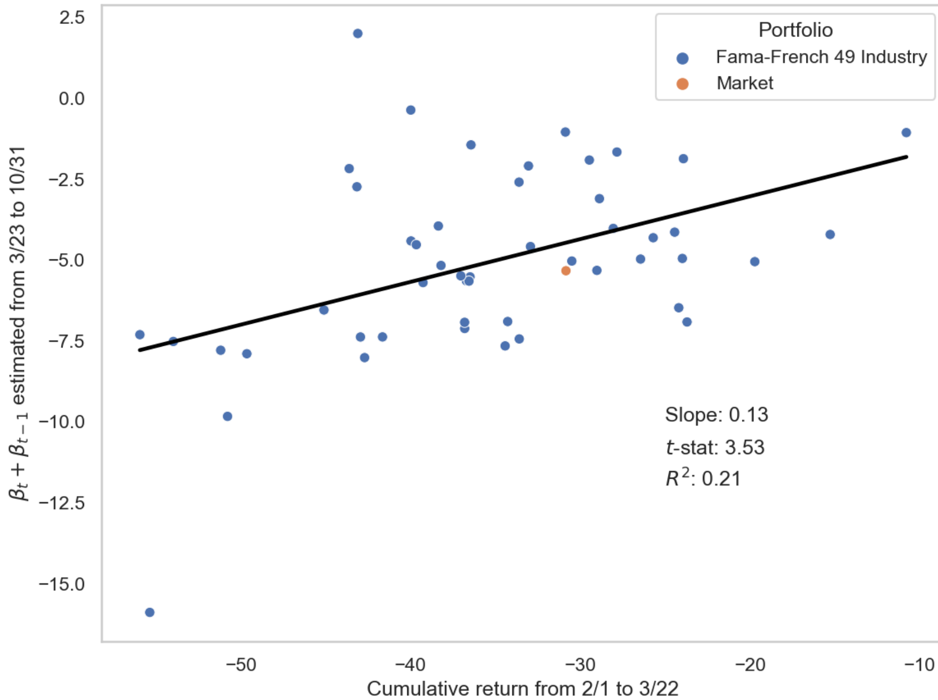

Note: Figure plots industry sensitivity to vaccine progress against exposure to COVID-19 as measured by cumulative returns. Cumulative returns are from February 1, 2020 to March 22, 2020. Sensitivity to vaccine progress is estimated from March 23, 2020 to October 31, 2020 as in (21).

sample from March 23, 2020 to October 31, 2020, by re-estimating (20) sector-by-sector,

$$R^e_{i,t} = \alpha + \sum_{h=-1}^{0} \beta_{h.i} \Delta \mathrm{VPI}_{t+h} + \gamma_{1,i} R^e_{i,t-1} + \gamma_2 R^e_{i,t-2} + \sum_{j=1}^{28} \delta_{j,i} \mathbb{1}_{\text{jump } j} + \epsilon_{i,t} \tag{21}$$

where $R^e_{i,t}$ is value-weight excess returns on the 49 Fama-French industry portfolios.

Figure 3 presents the results. Each industry's sensitivity to vaccine progress is plotted against its exposure to COVID-19. The relationship is negative and statistically significant – industries that were more exposed to COVID-19 subsequently saw more positive price impact as the vaccine was expected to deploy sooner. The industries also exhibit notable variation. Oil, fabricated products and recreation were among those with higher COVID-19 exposure and vaccine progress sensitivity, while pharmaceutical products, food products and computer software had lower exposure and sensitivity.

The association of industry exposure to COVID-19 with its subsequent sensitivity to our index lends confidence to the construction and interpretation of the index as, in fact, measuring vaccine progress.

4.2 Welfare Cost of the Pandemic

Referring now to Figure 2 in Section 2.6, we can see that, in the context of the model, stock market sensitivities in the range of 0.05 imply welfare costs per remaining expected year of the pandemic of roughly the same amount or slightly higher using the benchmark parameters from Table 1. As that section explained, this conclusion can be supported by any choice of the disaster parameters that matches the observed sensitivity. For example, the choices $(\mu(1) = 0.015, \sigma(1) = 0.075, \chi = 0.035, \zeta = 1)$ and $(\mu(1) = 0.03, \sigma(1) = 0.075, \chi = 0.025, \zeta = 2)$ each yields a sensitivity of 0.05 and each implies a welfare cost of 0.052 per year.

We can now combine this estimate with our forecasts for the expected time to successful vaccine deployment constructed in Section 3 (shown in Figure 1). Assuming that this series captures belief at the time for the pandemic's duration, and assuming investors' model of the pandemic economy aligned with the model's dynamics, our results imply that in late April 2020 when deployment was expected in about a year, the welfare value of ending the pandemic was approximately 5% of total wealth. At the beginning of March, with very little vaccine progress reported, the analogous numbers are approximately three years forecast duration, and a welfare cost of 15%. By November 2020, with less than six months expected until successful vaccine deployment, ending the pandemic immediately would still have been worth over 2% of total wealth. These are the paper's principal findings.

We defined V as the welfare gain from transitioning immediately to the non-pandemic state. More generally, the quantity (17) defined in Section 2.4 gives the the value of any intervention that shortens the expected duration of the crisis. For example, a partially successful vaccine technology which cuts this duration in half can be associated with its analogous welfare gain. Using one of the above choices for the disaster parameters, Table 5 shows the fraction of wealth the representative agent would be willing to pay to lower the expected duration from one value to another. The table entries are close to constant along each diagonal, indicating that the welfare gain scales almost linearly with the expected change in duration. A semi-effective vaccine that reduced the duration from four to two years is worth somewhat less but close to one that reduces the duration from two years to zero.

4.2.1 Discussion

In assessing our conclusions on the welfare cost of the pandemic, three natural questions arise. First, is 5-15% of total wealth a reasonable amount to be willing to pay to avoid or curtail a pandemic? Second, how should one think about "five percent of total wealth" in terms of real-world values (e.g., dollars)? Finally, how do these magnitudes compare to other ways of assessing the cost of COVID-19?

Within the context of the model, the welfare gain estimate is intimately tied to the rate of loss (or lower growth rate) of wealth due to the disaster shocks. In the calibrations reported above, parameters consistent with the observed stock market sensitivity all imply that the expected growth of wealth is 5%-6% less per year while the crisis is unchecked. Hence, sacrificing five percent of wealth per year of expected duration to avoid this outcome makes sense economically.

As another check, one could compare our estimate to the total drop in stock market wealth – around 36% – at the onset of the pandemic. This is not the same quantity that we are estimating. However, the two numbers are not unrelated. Any estimate of the cost of COVID-19 substantially less than 36% of stock market wealth would be difficult to reconcile with this observation.

Our estimate is also similar to magnitudes reported in the literature that computes the welfare gain of eliminating other types of disasters.[19] Barro (2009) reports that, in a model with rare disasters, moderate risk aversion, and an EIS greater than one, society would be willing to pay up to 20% of permanent income to eliminate disaster risk. Pindyck and Wang (2013) estimates the willingness to pay to reduce the impact of a disaster to 15% of capital stock at 7%.

Turning to the second question, as a baseline value, total U.S. household wealth at the end of 2019 was approximately \$96 trillion, implying that five percent represents about \$5 trillion. However, some ambiguity arises in interpretation because of the stylized nature of the model's depiction of wealth. A single state variable, q, represents not only household wealth, but the (book) value of the capital stock and the (market) value of a claim to all future dividends. One possibility is to view the magnitude of q through the lens of consumption. Aggregate U.S. nondurable and service consumption in 2019 was approximately \$13.4 trillion. In the model calibration, the marginal propensity to consume (C/q) is approximately 0.04,[20] which would imply that five percent of q represents $(0.05/0.04)$ times C, or $V \approx \$17$ trillion.

Dollar values in this range are plausible, and are not out of line with other estimates of the cost of the pandemic. There is now a substantial literature estimating the cost of COVID-19 based on foregone health and economic activity. As discussed in the introduction, a common approach in health economics is to assign values to lives and productivity lost due to the virus. Writing in mid-2020, Cutler and Summers (2020) forecast health losses (including both morbidity and mental health) caused by the pandemic and estimate the total economic cost of COVID-19 to be \$16 trillion under the assumption that "it will be substantially contained by the fall of 2021." Implicitly, then, this is an estimate of

[19] Values are commonly reported as percentage reductions of permanent income. Such numbers are directly comparable to our percentages of (permanent) reductions in q since consumption is proportional to q.

[20] This calibration aligns with estimates of Blundell et al. (2008) and Souleles (2002), among others.

a rate of loss for one year.[21] The similarity of this estimate to our own – despite the very different inputs and assumptions – lends credence to each.

4.2.2 Extensions

We extend the model in several directions in the online appendix.[22] A first extension endogenizes the real option to invest in vaccine research so that the speed of progress is an equilibrium outcome. Although we do not attempt to estimate a production function for pharmaceutical R&D, it is clear *a priori* that the more powerful the available technology the smaller the welfare cost of a pandemic. Nonetheless, we show that, given the observed market response to vaccine progress, and the observed expected duration of the pandemic, our welfare calculation would not be significantly altered under this version of the model.

A second generalization endogenizes the magnitude of health shocks via labor choice in order to study the role of exposure externalities. We assume labor augments agent's capital stock in production; however, it also exposes the agent to the pathogen. The agent thus optimally withdraws labor in the pandemic states, and the magnitude of the labor withdrawal then determines the equilibrium severity of the shocks to output. However, agents' privately optimal labor choice does not fully internalize the exposure created for other agents. Using our empirical estimates of the COVID-19 severity, and estimates in the literature of the withdrawal of labor in Spring 2020, we compute the difference in the welfare gain with individually optimal labor choice versus that by a central planner. Since the planner imposes a stricter labor contraction (or lockdown) and thus subjects the economy to less damage, the welfare gain is approximately 15% lower than under the

[21] Focusing just on GDP, CBO (2020) estimate over $7 trillion in lost output through 2030. The IMF's World Economic Outlook (IMF (2021)) estimate the collapse could have been three times as large had policymakers not enacted significant intervention (including $16 trillion in fiscal support). They further estimate the cumulative loss in output relative to the counterfactual without COVID-19 to be $28 trillion over 2020–2025.

[22] To our knowledge, these extensions have not previously been examined in the literature on the welfare cost of disasters or business cycle risk.

representative agent's policy. This difference rises with the severity of the externality as measured by the increased degree of lockdown under central planning.

Finally, the online appendix examines the extent to which the model can offer a plausible depiction of the path of the stock market and consumption that was actually experienced during 2020. For tractability, thee model omits many factors that played important roles in 2020. We have not included fiscal or monetary policy mechanisms, for example. Fully accounting for the behavior of the financial markets and the real economy during 2020 is beyond the scope of the paper. Although there are challenges, we argue that a consistent interpretation of the actual experience of 2020 is possible, and discuss the extent to which our conclusions may be affected by the model's limitations.

5 Conclusion

This paper provides an estimate of the value of reducing the expected duration of the COVID-19 pandemic using the joint behavior of stock prices and a novel vaccine progress indicator based on the chronology of stage-by-stage advance of individual vaccine candidates and related news during 2020. In the context of a general equilibrium regime-switching model of repeated pandemics, the sensitivity of the stock market to vaccine progress indicator is essentially determined by the expected rate of loss (or lower growth rate) of wealth during a pandemic. Our empirical estimate can thus be translated into an implied welfare gain attributable to reverting to the non-pandemic state. With standard preferences parameters, this gain is approximately 5-15% of wealth, depending on the expected remaining duration of the pandemic. This number can also be interpreted as a measure of the expected cost of the pandemic.

References

Hengjie Ai. Information quality and long-run risk: Asset pricing implications. *The Journal of Finance*, 65(4):1333–1367, 2010.

Fernando Alvarez and Urban J Jermann. Using asset prices to measure the cost of business cycles. *Journal of Political economy*, 112(6):1223–1256, 2004.

Asger Lau Andersen, Emil Toft Hansen, Niels Johannesen, and Adam Sheridan. Consumer responses to the covid-19 crisis: Evidence from bank account transaction data. Technical report, CEPR, April 2020.

Daniel Andrei, Michael Hasler, and Alexandre Jeanneret. Asset pricing with persistence risk. *The Review of Financial Studies*, 32(7):2809–2849, 11 2019.

Natalie Bachas, Peter Ganong, Pascal J Noel, Joseph S Vavra, Arlene Wong, Diana Farrell, and Fiona E Greig. Initial impacts of the pandemic on consumer behavior: Evidence from linked income, spending, and savings data. Working Paper 27617, National Bureau of Economic Research, July 2020.

Scott Baker, R.A. Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine Yannelis. How does household spending respond to an epidemic? consumption during the 2020 covid-19 pandemic. *NBER Working Paper*, 2020a.

Scott R Baker, Nicholas Bloom, Steven J Davis, Kyle Kost, Marco Sammon, and Tasaneeya Viratyosin. The Unprecedented Stock Market Reaction to COVID-19. *The Review of Asset Pricing Studies*, 07 2020b.

Robert J. Barro. Rare Disasters and Asset Markets in the Twentieth Century*. *The Quarterly Journal of Economics*, 121(3):823–866, 08 2006.

Robert J Barro. Rare disasters, asset prices, and welfare costs. *American Economic Review*, 99(1):243–64, 2009.

Richard Blundell, Luigi Pistaferri, and Ian Preston. Consumption inequality and partial insurance. *American Economic Review*, 98(5):1887–1921, December 2008.

CBO. An update to the economic outlook: 2020 to 2030. Technical report, 2020.

Haiqiang Chen, Welan Qian, and Qiang Wen. The impact of the covid-19 pandemic on consumption: Learning from high frequency transaction data. Technical report, Working Paper, 2020.

Raj Chetty, John N Friedman, Nathaniel Hendren, Michael Stepner, and The Opportunity Insights Team. How did covid-19 and stabilization policies affect spending and employment? a new real-time economic tracker based on private sector data. Working Paper 27431, National Bureau of Economic Research, June 2020.

Olivier Coibion, Yuriy Gorodnichenko, and Michael Weber. The cost of the covid-19 crisis: Lockdowns, macroeconomic expectations, and consumer spending. Working Paper 27141, National Bureau of Economic Research, May 2020.

Pierre Collin-Dufresne, Michael Johannes, and Lars A. Lochstoer. Parameter learning in general equilibrium: The asset pricing implications. *American Economic Review*, 106(3): 664–98, March 2016.

Matt Craven, Adam Sabow, Lieven Van der Venken, and Matt Wilson. Not the last pandemic: Investing now to reimagine public-health systems. Research report, McKinsey and Company, May 2021.

David M. Cutler and Lawrence H. Summers. The COVID-19 Pandemic and the $16 Trillion Virus. *JAMA*, 324(15):1495–1496, 10 2020.

Marco Del Angel, Caroline Fohlin, and Marc D Weidenmier. Do global pandemics matter for stock prices? lessons from the 1918 spanish flu. Technical report, National Bureau of Economic Research, January 2021.

Darrell Duffie and Larry G. Epstein. Asset pricing with stochastic differential utility. *Review of Financial Studies*, 5:411–436, 1992.

Darrell Duffie and Costis Skiadas. Continuous-time security pricing: A utility gradient approach. *Journal of Mathematical Economics*, 23:107–132, 1994.

Vadim Elenev, Tim Landvoigt, and Stijn Van Nieuwerburgh. Can the covid bailouts save the economy. Technical report, CEPR working Paper DP14714, 2020.

X. Gabaix. Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *Quarterly Journal of Economics*, 127(2):645–700, 2012.

Bill Gates. *How to Prevent the Next Pandemic*. Knopf, 2022.

François Gourio. Disaster risk and business cycles. *American Economic Review*, 102(6): 2734–66, May 2012.

Harrison Hong, Jeffrey Kubik, Neng Wang, Xiao Xu, and Jinqiang Yang. Reopening effect of COVID-19 vaccines on corporate earnings. Technical report, Working Paper, October 2021a.

Harrison Hong, Neng Wang, and Jinqiang Yang. Implications of stochastic transmission rates for managing pandemic risks. *Review of Financial Studies*, 34(11):5224–5265, 2021b.

Harrison Hong, Neng Wang, and Jinqiang Yang. Mitigating disaster risks in the age of climate change. Technical report, National Bureau of Economic Research, February 2022.

IMF. World economic outlook. Technical report, 2021.

Òscar Jordà, Moritz Schularick, and Alan M Taylor. Disasters everywhere: The costs of business cycles reconsidered. Working Paper 26962, National Bureau of Economic Research, August 2020.

Eric S. Lander and Jacob J. Sullivan. American pandemic preparedness: Transforming our capabilities. White paper, White House, September 2021.

Robert Lucas. *Models of Business Cycles*. Blackwell, Oxford, 1987.

Ian WR Martin. Disasters and the welfare cost of uncertainty. *American Economic Review*, 98(2):74–78, 2008.

Ian WR Martin and Robert S Pindyck. Averting catastrophes: the strange economics of scylla and charybdis. *American Economic Review*, 105(10):2947–85, 2015.

Ian WR Martin and Robert S Pindyck. Welfare costs of catastrophes: lost consumption and lost lives. *The Economic Journal*, 131(634):946–969, 2021.

Emi Nakamura, Jón Steinsson, Robert Barro, and José Ursúa. Crises and recoveries in an empirical model of consumption disasters. *American Economic Journal: Macroeconomics*, 5(3):35–74, 2013.

Robert S. Pindyck and Neng Wang. The economic and policy consequences of catastrophes. *American Economic Journal: Economic Policy*, 5(4):306–39, November 2013.

Esther S Pronker, Tamar C Weenen, Harry Commandeur, Eric HJHM Claassen, and Albertus DME Osterhaus. Risk in vaccine research and development quantified. *PloS one*, 8(3):e57755, 2013.

Adam Sheridan, Asger Lau Andersen, Emil Toft Hansen, and Niels Johannesen. Social distancing laws cause only small losses of economic activity during the covid-19 pandemic in scandinavia. *Proceedings of the National Academy of Sciences*, 117(34):20468–20473, 2020.

Nicholas S Souleles. Consumer response to the Reagan tax cuts. *Journal of Public Economics*, 85(1):99–120, 2002.

Lawrence Summers, Ngozi Okonjo-Iweala, and Tharman Shanmugaratnam. Report of the G20 high level independent panel on financing the global commons for pandemic preparedness and response. Report, G20, September 2021. URL https://pandemic-financing.org/report.

Thomas D Tallarini Jr. Risk-sensitive real business cycles. *Journal of Monetary Economics*, 45(3):507–532, 2000.

Jerry Tsai and Jessica A. Wachter. Disaster risk and its implications for asset pricing. *Annual Review of Financial Economics*, 7(1):219–252, 2015.

Chi Heem Wong, Kien Wei Siah, and Andrew W Lo. Estimation of clinical trial success rates and related parameters. *Biostatistics*, 20(2):273–286, 01 2018.

Table 1: Parameter Values

Parameter	Symbol	Value
	Preferences	
Coefficient of relative risk aversion	γ	4.0
Elasticity of intertemporal substitution	ψ	1.5
Rate of time preference	ρ	0.0125
	Normal Output	
Expected growth	$\mu(0)$	0.03
Volatility	$\sigma(0)$	0.05
	Disaster Timing	
Disaster intensity	η	0.02
Disaster duration	$1/\lambda$	4.0

The table shows the baseline parameter values used in inferring the welfare cost of the disaster, V, from the stock market reaction to disaster duration news.

Table 2: Forecast Comparison

Date	Survey median	VPI	% respondents below
	Deutsche Bank		
Date	Survey median	VPI	% respondents below
May	1.158	0.958	35.0
June	1.162	0.893	31.2
July	0.920	0.595	20.8
Sep	0.625	0.561	44.3
	Superforecasters		
Date	Survey median	VPI	% respondents below
April	1.902	1.291	16.1
May	1.603	0.958	14.6
June	1.189	0.893	31.0
July	0.808	0.595	32.7
August	0.519	0.606	58.4
September	0.445	0.518	57.2

Note: The table compares forecasts for the earliest date of vaccine availability in years. The top panel compares the median from a survey conducted by Deutsche Bank, while the bottom panel compares the median from a survey conducted by Good Judgement Inc. The column VPI denotes the forecast from our estimated vaccine progress indicator, and the last column reports the percent of respondents from each survey with forecasts below ours. Survey respondents are reported in calendar intervals. The comparison assumes a uniform distribution of forecasts in time within the median bin. The survey dates are as of the end of the month in the first column, except the Deutsche Bank September survey which is for the week ending September 11, 2020.

Table 3: Stock Market Sensitivity to Vaccine Progress

	(1)	(2)	(3)	(4)	(5)	(6)
γ_1	-0.091	-0.068	-0.082	-0.063	-0.098	-0.089
	(0.068)	(0.066)	(0.066)	(0.064)	(0.069)	(0.069)
γ_2	0.150*	0.180*	0.154*	0.177*	0.138	0.150*
	(0.090)	(0.091)	(0.092)	(0.090)	(0.086)	(0.083)
β_t	-1.229	-2.084*	-1.409	-2.069*	-1.000	-1.333
	(1.353)	(1.193)	(1.290)	(1.243)	(1.141)	(1.114)
β_{t-1}			-3.836	-4.523	-3.850	-4.269
			(3.069)	(3.093)	(2.892)	(2.849)
$\beta_{t-5,t-2}$					-1.959***	-2.281***
					(0.689)	(0.631)
α	0.290***		0.225**		0.149	
	(0.105)		(0.097)		(0.099)	
$\sum_h \beta_{t-h}$	-1.229	-2.084	-5.246	-6.592	-6.808	-7.883
F-stat on $\sum_h \beta_{t-h}$	0.83	3.05	4.51	7.20	7.01	10.54
P-value on $\sum_h \beta_{t-h}$	0.36	0.08	0.04	0.01	0.01	0.00
N	203	203	203	203	203	203

Note: The table shows results from regression (20). The dependent variable is daily excess returns on the market portfolio in percent. The return on the value-weighted CRSP index is used from January 1, 2020 to October 31, 2020. Independent variables include two lags of excess returns on the market portfolio, changes in vaccine progress indicator in years, and dummy variables for each jump date from Baker et al. (2020b) unrelated to news about vaccine progress. All columns are employ the baseline construction of the vaccine forecast. Newey-West standard errors with four lags are shown in parentheses.Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Stock Market Sensitivity to Vaccine Progress – Robustness

	(1)	(2)	(3)	(4)	(5)	(6)
News	All states	None	Current state	All states	All states	All states
Depreciation	Y	N	Y	Y	Y	Y
Cor(n, n')	0.2	0.2	0.2	0.4	0.2	0.2
$\pi^{\text{base}}_{\text{approval}}$	0.95	0.95	0.95	0.95	0.85	0.95
Ex-China and Russia	Y	Y	Y	Y	Y	N
γ_1	-0.082	-0.081	-0.083	-0.086	-0.083	-0.091
	(0.066)	(0.065)	(0.065)	(0.066)	(0.066)	(0.056)
γ_2	0.154*	0.154*	0.153*	0.153*	0.154*	0.143
	(0.091)	(0.092)	(0.092)	(0.091)	(0.091)	(0.087)
β_t	-1.403	-1.829	-1.605	-1.577	-1.471	1.073
	(1.287)	(2.055)	(1.338)	(1.232)	(1.226)	(1.676)
β_{t-1}	-3.829	-4.715	-2.725	-3.308	-3.691	-5.812**
	(3.063)	(3.897)	(2.595)	(3.092)	(2.953)	(2.921)
α	0.227**	0.207*	0.243**	0.236**	0.228**	0.228**
	(0.095)	(0.098)	(0.096)	(0.095)	(0.095)	(0.095)
$\sum_h \beta_{t-h}$	-5.233	-6.544	-4.330	-4.994	-5.161	-4.739
F-stat on $\sum_h \beta_{t-h}$	4.51	5.59	4.93	3.73	4.49	3.71
P-value on $\sum_h \beta_{t-h}$	0.04	0.02	0.03	0.06	0.04	0.06
N	206	206	206	206	206	206

Note: Table shows the results from specification (1) in the paper. The dependent variable is daily percent excess returns on the market portfolio. Independent variables include two lags of excess returns on the market portfolio, a five-day window of changes in vaccine progress indicator in years, and dummy variables for each jump date from Baker et al. (2020b) unrelated to news about vaccine progress. The first column is the baseline specification with news applying to all states, deterministic depreciation, base copula correlation of 0.2, probability of success in the application state equal to 0.95 and excludes candidates from China and Russia. Column 2 removes news and depreciation; 3 restricts news to the current state and increases the $\Delta\pi$ from news on positive data releases, positive enrollment and dose starts to 15%, 5% and 5%, respectively; 4 doubles the base copula correlation to 0.4; 5 decreases the probability of success to 0.85 in the application state; and 6 includes candidates from China and Russia. The return on the value-weighted CRSP index is used from January 1, 2020 to October 31, 2020. The table uses Newey-West standard errors with 4 lags (shown in parentheses). Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Welfare Gain as a Function of Reduction in Expected Duration

	$T2$: EXPECTED DURATION AFTER INTERVENTION (YEARS)							
$T1$: INITIAL EXPECTED DURATION	3.50	3.00	2.50	2.00	1.50	1.00	0.50	0
4.00	2.60	5.15	7.65	10.09	12.48	14.82	17.11	19.35
3.50	--	2.62	5.18	7.69	10.14	12.54	14.89	17.19
3.00	--	--	2.63	5.21	7.73	10.19	12.61	14.97
2.50	--	--	--	2.64	5.23	7.77	10.24	12.67
2.00	--	--	--	--	2.66	5.26	7.81	10.29
1.50	--	--	--	--	--	2.67	5.29	7.84
1.00	--	--	--	--	--	--	2.69	5.31
0.50	--	--	--	--	--	--	--	2.70

Note: The table shows the percentage of wealth that the representative would be willing to trade for an intervention that shortens the pandemic from an initial expected duration of $T1$ years to another state with $T2 < T1$ years remaining in expectation. The pandemic parameters are $\mu(1) = 0.015, \sigma(1) = 0.075, \chi = 0.035$, and $\zeta = 1$. The remaining parameters are shown in Table 1.