Internet Appendix

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

February 2023

Appendix A contains proofs of the propositions in Section 2 of the paper, and describes a decentralization of the economy that motivates viewing the output claim as representing the market portfolio.

Appendix B includes additional details on construction of the vaccine progress indicator as described in Section 3.

Appendix C presents extensions of the model to include endogenous investment in a mitigation technology, and endogenous labor supply.

Appendix D discusses the experiences of the stock market and consumption during 2020 as interpreted through the lense of the model.

A Proofs

A.1 Proof of Proposition 1

Proposition 1. *Denote*

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

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

$$g_0 \equiv g(0) = \frac{(1-\gamma)}{(\psi-1)}\rho^{\psi}\left(H(0)\right)^{-\psi\theta^{-1}} + \eta\left[\frac{H(1)}{H(0)} - 1\right] \tag{A.2}$$

$$g_1 \equiv g(1) = \frac{(1-\gamma)}{(\psi-1)}\rho^{\psi}\left(H(s)\right)^{-\psi\theta^{-1}} + \lambda_d\left[\frac{H(s-1)}{H(s)} - 1\right] + \lambda_u\left[\frac{H(s+1)}{H(s)} - 1\right], \tag{A.3}$$

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

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

$$C(s) = \frac{\left(H(s)\right)^{-\psi\theta^{-1}} q}{\rho^{-\psi}}, \tag{A.4}$$

and the value function of the representative agent is

$$\mathbb{J}(s) \equiv \frac{H(s) q^{1-\gamma}}{1-\gamma}. \tag{A.5}$$

Proof. From the evolution of capital stock for the representative agent (8), we obtain the Hamilton-Jacobi-Bellman (HJB) equation as follows for each state s:

$$\begin{aligned} 0 = \max_C \Big[& f(C, \mathbb{J}(s)) - \rho\mathbb{J}(s) + \mathbb{J}_q(s)(q\mu(s) - C) \\ & + \frac{1}{2}\mathbb{J}_{qq}(s) q^2 \sigma(s)^2 + \zeta(s)\left[\mathbb{J}(s)\left(q(1-\chi(s))\right) - \mathbb{J}(s)(q)\right] \\ & + \lambda_u(s)\left[\mathbb{J}(s+1)(q) - \mathbb{J}(s)(q)\right] + \lambda_d(s)\left[\mathbb{J}(s-1)(q) - \mathbb{J}(s)(q)\right] \Big] \end{aligned} \tag{A.6}$$

Taking the first-order condition with respect to $C(s)$ in HJB equation (A.6), we obtain

$$f_c(C,\mathbb{J}(s)) - \mathbb{J}_q(s) = 0. \tag{A.7}$$

Using $f(C,\mathbb{J})$ from (10) and taking the derivative with respect to C, we obtain

$$f_c = \frac{\rho C^{-\psi^{-1}}}{[(1-\gamma)\mathbb{J}(s)]^{\frac{1}{\theta}-1}}. \tag{A.8}$$

Substituting the conjecture $\mathbb{J}(s)$ in equation (A.5) yields

$$f_c = \frac{\rho C^{-\psi^{-1}}}{H(s)^{\frac{\gamma-\psi^{-1}}{1-\gamma}} q^{\gamma-\psi^{-1}}}. \tag{A.9}$$

Then, for state $s \in \{0,\ldots,S\}$, we obtain by substituting $\mathbb{J}_q(s) = H(s)q^{-\gamma}$ in (A.7), and simplifying:

$$C(s) = \frac{H(s)^{-\theta\psi^{-1}} q}{\rho^{-\psi}}. \tag{A.10}$$

To verify the conjectured form of the value function, we plug it in to the HJB equation (A.6) and reduce it to the recursive system in the proposition via the following steps:

1. substitute the optimal policy $C(s)$ into the HJB equation (A.6);
2. cancel the terms in q which have the same exponent; and
3. group terms not involving $H(s)$ constants into $g(0)$ for state $s = 0$ and $g(s)$ for state $s \in \{1,\ldots,S-1\}$

to reach equations (A.1) – (A.3). This system of recursive equations can be solved numerically with the final condition in Proposition 2: $H(s) = H(0)$, that states 0 and S are non-disaster states. □

□

A.2 Proof of Proposition 2

Proposition 2. *The welfare gain to ending the disaster 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):*

$$V(s) = 1 - \left(\frac{c(s)}{c(0)}\right)^{-\frac{1}{\psi-1}} = 1 - \left(\frac{C(s)}{C(0)}\right)^{-\frac{1}{\psi-1}} \tag{A.11}$$

Proof. The welfare gain to ending the disaster $V(s)$ satisfies:

$$\mathbb{J}(0)(q) = \mathbb{J}(0)\left[(1 - V(s))\, q\right] \tag{A.12}$$

where $\mathbb{J}(0)$ is evaluated at $(1 - V(s))\, q$. Substituting for $\mathbb{J}(s)$ from (A.5), we obtain

$$\frac{H(0)q^{1-\gamma}}{(1-\gamma)} = \frac{H(0)\left[(1 - V(s))\, q\right]^{1-\gamma}}{(1-\gamma)} \tag{A.13}$$

which yields

$$V(s) = 1 - \left(\frac{H(s)}{H(0)}\right)^{\frac{1}{1-\gamma}}. \tag{A.14}$$

Then, substituting for $C(s)$ from (A.4) and recognizing marginal propensity to consume, $c(s)$, equals $\frac{dC}{dq} = \frac{C(s)}{q}$, yields Proposition 2. $\square$

A.3 Proof of Proposition 3

This proof of Proposition 3 treats the case of the generalized version of the model with endogenous labor supply.

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.*

Proof. To begin, we derive the pricing kernel and the risk-free rate. Under stochastic differential utility, the kernel can be represented as

$$\Lambda_t = e^{\int_0^t f_{\mathbb{J}} du} f_C \tag{A.15}$$

where

$$f(C, J) = \rho \frac{C^{\varrho}}{\varrho} \left((1-\gamma)\mathbb{J}\right)^{1-\frac{1}{\theta}} - \rho\theta\mathbb{J} \tag{A.16}$$

where $\varrho = 1 - \frac{1}{\psi}, \theta = \frac{1-\gamma}{\varrho}$. As shown in Section 4, the value function and the consumption flow rates are:

$$\mathbb{J} = q^{1-\gamma} H(s)/(1-\gamma) \quad \text{and} \quad C = \rho^{\psi} H(s)^{e}\, q(s) q \tag{A.17}$$

where $e = \frac{1-\psi}{1-\gamma}$. Together these imply

$$f_C = \rho C^{\varrho-1} \left((1-\gamma)\mathbb{J}\right)^{1-\frac{1}{\theta}} \tag{A.18}$$

or

$$f_C = \rho \left(\rho^{\psi} H(s)^{e} q\right)^{\varrho-1} \left((1-\gamma)\left(q^{1-\gamma} H(s)/(1-\gamma)\right)\right)^{1-\frac{1}{\theta}}. \tag{A.19}$$

Simplifying, we get:

$$f_C = \rho^{1+\psi(\varrho-1)} H(s)^{e(\varrho-1)+\frac{\theta-1}{\theta}} q^{(\varrho-1)+\frac{(1-\gamma)(\theta-1)}{\theta}}. \tag{A.20}$$

The exponent of ρ is: $1+\psi(\varrho-1) = 1+\psi(-\frac{1}{\psi}) = 0$. The exponent of q is: $(\varrho-1)+\frac{(1-\gamma)(\theta-1)}{\theta}$. Substitute $\theta = \frac{1-\gamma}{\varrho}$ to get: $(\varrho-1)+\varrho(\frac{1-\gamma}{\varrho}-1) = -\gamma$. The exponent of $H(s)$ is

$$e(\varrho-1) + \frac{\theta-1}{\theta} \Rightarrow \frac{1-\psi}{1-\gamma}\left(-\frac{1}{\psi}\right) + \frac{1-\gamma\psi}{\psi(1-\gamma)} = 1 \tag{A.21}$$

Hence, $f_C = H(s)q^{-\gamma}$. Next, to evaluate $f_{\mathbb{J}}$, note that

$$f_{\mathbb{J}} = \rho \frac{C^{\varrho}}{\varrho} \left(1 - \frac{1}{\theta}\right) \left[(1-\gamma)\mathbb{J}\right]^{-\frac{1}{\theta}} (1-\gamma) - \rho\theta \tag{A.22}$$

Plugging in for C and $\mathbb{J}$ we get:

$$f_{\mathbb{J}} = \rho \frac{\left(\rho^{\psi} H(s)^{e} q\right)^{\varrho}}{\varrho} \left(1 - \frac{1}{\theta}\right) \left[(1-\gamma)\left(q^{1-\gamma} H(s)/(1-\gamma)\right)\right]^{-\frac{1}{\theta}} (1-\gamma) - \rho\theta \tag{A.23}$$

or

$$f_{\mathbb{J}} = \rho \frac{\left(\rho^{\psi} H(s)^{e} q\right)^{\varrho}}{\varrho} \left(\frac{\theta - 1}{\theta}\right) \left[\left(q^{1-\gamma} H(s)\right)\right]^{-\frac{1}{\theta}} (1-\gamma) - \rho\theta. \tag{A.24}$$

This can be expressed as:

$$f_{\mathbb{J}} = \frac{1}{\varrho} \rho^{1+\psi\varrho} H(s)^{e\varrho} q^{\varrho} \left(\frac{\theta - 1}{\theta}\right) (1-\gamma) q^{\frac{\gamma-1}{\theta}} H(s)^{-\frac{1}{\theta}} - \rho\theta. \tag{A.25}$$

Collecting terms:

$$f_{\mathbb{J}} = \frac{1}{\varrho} \rho^{1+\psi\varrho} H(s)^{e\varrho - \frac{1}{\theta}} q^{\varrho + \frac{\gamma-1}{\theta}} \left(\frac{\theta - 1}{\theta}\right) (1-\gamma) - \rho\theta. \tag{A.26}$$

Here the exponent of ρ is : $1 + \psi\varrho = \psi$, and the exponent of $H(s)$ is: $e\varrho - \frac{1}{\theta} = e\varrho - \frac{\varrho}{1-\gamma} = e$, and the exponent of q is: $\varrho + \frac{\gamma-1}{\theta} = 0$. Hence,

$$f_{\mathbb{J}} = \frac{1}{\varrho} \rho^{\psi} H(s)^{e} \left(\frac{\theta - 1}{\theta}\right) (1-\gamma) - \rho\theta = \rho^{\psi} H(s)^{e} (\theta - 1) - \rho\theta = c(s)(\theta - 1) - \rho\theta. \tag{A.27}$$

So, we conclude that

$$\Lambda_t = e^{\int_0^t f_{\mathbb{J}} du} f_C = q^{-\gamma} H(s) e^{\int_0^t [c(s)(\theta-1) - \rho\theta] du}. \tag{A.28}$$

The riskless interest rate, $r(s)$ is minus the expected change of $d\Lambda/\Lambda$ per unit time. Applying Itô's lemma to the above expression yields drift (or dt terms)

$$c\ (\theta - 1) - \rho\theta - \gamma(\ell^{\alpha}\mu - c) + \gamma(\gamma + 1)\ell^{\alpha}\sigma^{2} \tag{A.29}$$

where $\ell(0) = \bar{\ell} = 1$ and $\ell(s) = \ell^{\star}$ for $s > 0$. Note that the term $(\ell^{\alpha}\mu - c)$ is the drift of dq/q. To these terms we add the expected change from the jumps in the state s for $s = 0$:

$$\eta \left(\frac{H(1)}{H(0)} - 1 \right) \equiv \tilde{\eta} - \eta \tag{A.30}$$

which serves to define the risk-neutral jump intenstity $\tilde{\eta}$. For $s > 0$ the expected jumps include both up and down changes in s as well as jumps in $q^{-\gamma}$:

$$\lambda_u \left(\frac{H(s+1)}{H(s)} - 1 \right) + \lambda_d \left(\frac{H(s-1)}{H(s)} - 1 \right) + \zeta((1-\chi)^{-\gamma} - 1) \equiv (\tilde{\lambda_u} - \lambda_u) + (\tilde{\lambda_d} - \lambda_d) + (\tilde{\zeta} - \zeta) \tag{A.31}$$

where the risk neutral intensities are defined as for η. The full expression for $r(0)$ is then

$$-\left\{ c(0)\ (\theta - 1) - \rho\theta - \gamma(\mu - c(0)) + \gamma(\gamma+1)\sigma^2 + (\tilde{\eta} - \eta) \right\}. \tag{A.32}$$

For $s > 0$ we have $r(s)$ as

$$-\left\{ c(s)(\theta - 1) - \rho\theta - \gamma((\ell^\star)^\alpha \mu - c(s)) + \frac{1}{2}\gamma(\gamma+1)(\ell^\star)^\alpha \sigma^2 + (\tilde{\lambda_u} - \lambda_u) + (\tilde{\lambda_d} - \lambda_d) + (\tilde{\zeta} - \zeta)) \right\}. \tag{A.33}$$

We return to these expressions after deriving the pricing equation for the output claim.

By the fundamental theorem of asset pricing, the instantaneous expected excess return to the claim $P(q,s)$ must equal minus covariance of the returns to P with the pricing kernel. Deriving these two quantities and setting them equal yields the pricing system, to which the proof will construct the solution.

The expected excess return to the claim $P(q,s)$ is the sum of its expected capital gain and its expected payout, minus rP. In the nondisaster state, this is

$$\frac{1}{2}\sigma^2 q^2 P_{qq}(q,0) + (\mu - c(0))qP_q(q,0) + \eta(P(q,1) - P(q,0)) + \mu q - r(0)P(q,0) \tag{A.34}$$

whereas in the disaster states it is

$$\begin{aligned}
&\frac{1}{2}(\ell^\star)^\alpha\sigma^2q^2P_{qq}(q,s) + ((\ell^\star)^\alpha\mu - c(s))qP_q(q,s) \\
&+\lambda_u(P(q,s+1) - P(q,s)) + \lambda_d(P(q,s-1) - P(q,s)) + \zeta(P((1-\chi)q,s) - P(q,s)) \\
&+\mu(\ell^\star)^\alpha q - \zeta\chi q - r(s)P(q,s).
\end{aligned} \tag{A.35}$$

Next, we need to derive the covariance of the returns to P with $d\Lambda/\Lambda$. As mentioned in the text, in addition to the usual contribution of covariance from the capital gains dP/P, the covariance also includes the contribution from the dividends themselves, which are risky in this model. There are also contributions from both Brownian comovement and co-jumps in q and s. The Brownian terms are

$$-\gamma(\ell^\star)^\alpha\sigma^2[qP_q(q,s) + q] \tag{A.36}$$

for $s > 0$, or just $-\gamma\sigma^2[qP_q + q]$ for $s = 0$. The co-jump terms for $s > 0$ are

$$\begin{aligned}
&\zeta[P((1-\chi)q,s) - P(q,s) - \chi q]\ [(1-\chi)^{-\gamma} - 1] \\
&+\lambda_u[P(q,s+1) - P(q,s)]\left[\frac{H(s+1)}{H(s)} - 1\right] + \lambda_d[P(q,s-1) - P(q,s)]\left[\frac{H(s-1)}{H(s)} - 1\right]
\end{aligned} \tag{A.37}$$

or

$$\begin{aligned}
&[P((1-\chi)q,s) - P(q,s) - \chi q]\ [\tilde{\zeta} - \zeta] \\
&+[P(q,s+1) - P(q,s)][\tilde{\lambda_u} - \lambda_u] + [P(q,s-1) - P(q,s)][\tilde{\lambda_d} - \lambda_d].
\end{aligned} \tag{A.38}$$

For $s = 0$ the corresponding expression is just

$$[P(q,1) - P(q,0)][\tilde{\eta} - \eta]. \tag{A.39}$$

We now equate the expected excess return to minus the above covariance to obtain the difference/differential equation system that P must solve. Rather than repeating the general expressions, we instead conjecture that the the solutions are linear in q and deduce the resulting system. Under linearity $P_{qq} = 0$ and $P_q = p$, a constant that depends on s.

Plugging in the conjectured form, and cancelling a q, in states $s > 0$ the pricing equation says

$$((\ell^\star)^\alpha \mu - c(s))p(s) + \lambda_u(p(s+1) - p(s)) + \lambda_d(p(s-1) - p(s)) \tag{A.40}$$
$$- \chi\zeta p(s) + \mu(\ell^\star)^\alpha - \zeta\chi - r(s)p(s) - \gamma(\ell^\star)^\alpha\sigma^2[p(s)+1] \tag{A.41}$$
$$- \chi[p(s)+1]\ [\tilde{\zeta} - \zeta] + [p(s+1) - p(s)][\tilde{\lambda_u} - \lambda_u] + [p(s-1) - p(s)][\tilde{\lambda_d} - \lambda_d]$$
$$= 0. \tag{A.42}$$

Leaving the constant terms on the left, the right side consists of

$$p(s+1) \quad \text{terms:} \quad -\lambda_u - [\tilde{\lambda_u} - \lambda_u] = -\tilde{\lambda_u}, \tag{A.43}$$

$$p(s-1) \quad \text{terms:} \quad -\lambda_d - [\tilde{\lambda_d} - \lambda_d] = -\tilde{\lambda_d}, \tag{A.44}$$

and $p(s)$ terms:

$$-((\ell^\star)^\alpha\mu - c(s)) + \lambda_u + \lambda_d + \chi\zeta + r(s) + \gamma(\ell^\star)^\alpha\sigma^2 + \chi[\tilde{\zeta} - \zeta] + [\tilde{\lambda_u} - \lambda_u] + [\tilde{\lambda_d} - \lambda_d] \tag{A.45}$$

or

$$r(s) + c(s) - (\ell^\star)^\alpha(\mu - \gamma\sigma^2) + \tilde{\lambda_u} + \tilde{\lambda_d} + \chi\tilde{\zeta}. \tag{A.46}$$

The remaining constants on the left are

$$\mu(\ell^\star)^\alpha - \zeta\chi - \gamma(\ell^\star)^\alpha\sigma^2 - \chi[\tilde{\zeta} - \zeta]. \tag{A.47}$$

or

$$(\ell^\star)^\alpha(\mu - \gamma\sigma^2) - \chi\tilde{\zeta}. \tag{A.48}$$

The above equations define a linear system for $p(1)$ to $p(S-1)$. The pricing equation for $s = 0$ says

$$(\mu - c(s))p(0) + \eta(p(1) - p(0)) + \mu - r(0)p(0) - \gamma\sigma^2[p(0)+1] + [p(1) - p(0)][\tilde{\eta} - \eta] = 0, \tag{A.49}$$

or

$$\mu - \gamma\sigma^2 = p(0)[r(0) + c(0) - (\mu - \gamma\sigma^2) + \tilde{\eta}] - p(1)\,\tilde{\eta}. \tag{A.50}$$

This equation closes the system on the low end. At the high end, the system is closed via $p(S) = p(0)$.

Altogether the system may be written in matrix form,

$$\begin{bmatrix} r(0) + c(0) - (\mu - \gamma\sigma^2) + \tilde{\eta} & -\tilde{\eta} & 0 & \cdots \\ -\tilde{\lambda}_d & r(s) + c(s) - (\ell^*)^\alpha(\mu - \gamma\sigma^2) + \chi\tilde{\zeta} + \tilde{\lambda}_d + \tilde{\lambda}_u & -\tilde{\lambda}_u & 0 \\ 0 & \ddots & \ddots & \ddots \\ \vdots & \ddots & \ddots & \ddots \\ -\tilde{\lambda}_u & 0 & \cdots & \end{bmatrix} p = \begin{bmatrix} (\mu - \gamma\sigma^2) \\ (\ell^\star)^\alpha(\mu - \gamma\sigma^2) - \chi\tilde{\zeta} \\ \vdots \\ \vdots \\ \vdots \end{bmatrix}.$$

Assuming the parameters are such that the right-hand matrix is of full rank, the system has a unique, finite solution.

□

A.4 Output Claim

The model in Section 2 views the stock market as a claim to the economy's future output flow, defined as the change in the stock of wealth before consumption. This subsection clarifies the reasons for this definition, offers a decentralization that supports it, and also notes alternative claims that can achieve the same objectives.

First, our aim is to tractably depict a market that responds to news about vaccine progress, as captured by the state variable s. As is well known, in an economy where the capital stock can be costlessly converted to consumption goods, the unit price of the capital stock is constant and equal to 1.0. Or, in our notation, the value of q is q. Moreover, this is also equal to the value of a claim to all future consumption. In many applications, this is still a reasonable depiction of the market portfolio. But it does not describe the dynamic that we are interested in modeling.[1]

The consumption claim is a knife-edge case, however. Any wedge between consumption and

[1]Recall that in our model changes in the state s do not directly alter q.

payouts to equity will result in a nonconstant ratio P/q. In particular, claims whose cash flows vary more with the state of the pandemic will have prices that do so as well. The key feature of the claim that we price is that its expected cash flow mirrors the impact on wealth of the pandemic.

To see how the flow that we are valuing could constitute the payouts to owners of corporate claims, consider the following decentralization.

1. Households own the capital stock and rent it to goods-producing firms.
2. Firms produce output $\mu(s)\, q\, dt + \sigma(s)\, q\, dW_t$ per unit time.
3. Firms purchase insurance against pandemic shocks $-\chi q\, dJ$ from an insurance sector.
4. The market portfolio consists of a claim to the profits of both sectors plus the rental contract for the capital stock.

The assumption that firms rent productive capital from households is standard in macroeconomics. Notice that in step 1, the rental is effectively a riskless bond in that the "face value" of q is insured. Thus in this economy households separate risky and safe claims.

We assume the parameter values are such that positivity in all states holds, and we verify this for each case in our numerical work. We also verify that $P/q < 1$, i.e., that financial claims in the portfolio comprise a subset of the economy's total wealth.

Finally, it is worth clarifying that it is not necessary for our results to assume that holders of the portfolio bear the losses of the pandemic shock as a negative dividend. Our results are mathematically the same for an alternative claim that pays the (risk-neutral) expected output rate per unit time in each state.

B Vaccine Progress Indicator

This section includes additional details on construction of the vaccine progress indicator. Each day's forecast is computed via simulation as described in Section 2 of the paper. The procedure is depicted graphically in the flow chart Figure A.1.

The simulation takes as input a timeline of COVID-19 vaccine candidates' stage-by-stage progress from the London School of Hygiene & Tropical Medicine.[2] We observe the start dates

[2]This version of the paper uses the timeline available on November 2, 2020.

Figure A.1: Simulation Flow Chart

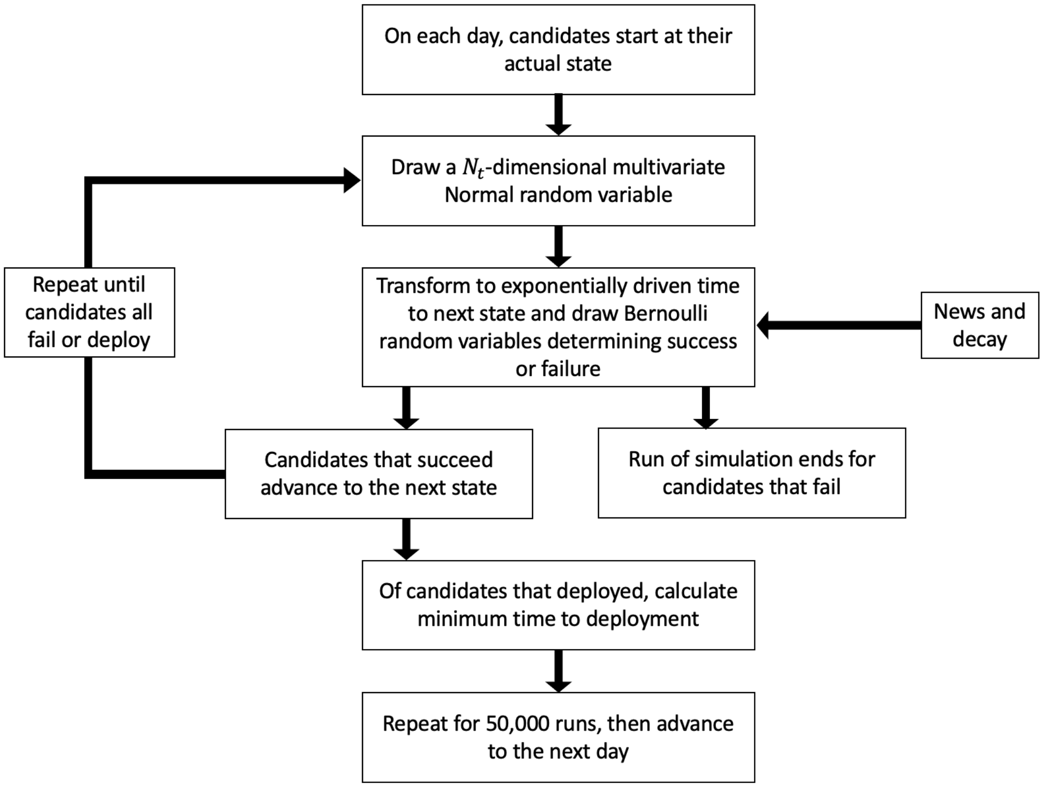

Note: Figure shows the simulation that estimates the expected time until vaccine deployment.

of each pre-clinical and clinical trial, along with their vaccine strategy. Table A.1 breaks down the number of candidates at each state at the end of our sample. We also observe each candidate's strategy. Table A.2 summaries the main strategies along with the number of candidates following each.

We then augment π_s^{base} with 233 news articles from FactSet StreetAccount, split into positive and negative news types. Table A.3 shows the number of articles by news type, while Table A.4 shows the top ten candidates by news count.

B.1 Data and Parameters

The simulation takes as input a timeline of COVID-19 vaccine candidates' stage-by-stage progress from the London School of Hygiene & Tropical Medicine. We observe the start dates of each pre-clinical and clinical trial, along with their vaccine strategy. Vaccines typically take years to

Table A.1: Vaccine States

State	# Candidates	Example Candidates
Preclinical	210	Amyris Inc Baylor College of Medicine Mount Sinai
Phase I Safety Trials	20	Clover/GSK/Dynavax CSL/University of Queensland Imperial College London
Phase II Expanded Trials	18	Arcturus/Duke Osaka/AnGes/Takara Bio Sanofi Pasteur/GSK
Phase III Efficacy Trials	11	AstraZeneca/Oxford BioNTech/Fosun/Pfizer Moderna

Note: Table describes the number of vaccine candidates in each state, along with example institutes. Data are from the London School of Hygiene & Tropical Medicine's COVID-19 Tracker. Data are as of November 2, 2020.

develop, and institutes have combined phases in an effort to accelerate the timeline. Following Wong et al. (2018), we adopt each candidate's most advanced state. We also observe each candidate's strategy.

Since candidates share a common virus target, and potentially common institutes or strategies, we define pairwise correlations in an additive manner. For two candidates $n \neq n'$:

$$\rho(n, n') = \begin{cases} 0.2 & \text{baseline} \\ \text{add } 0.2 & \text{if shared institute} \\ \text{add } 0.1 & \text{if shared strategy.} \end{cases}$$

Table A.5 lists our parameter choices of durations and baseline probabilities of success. Our baseline success probabilities are based upon estimates in Pronker et al. (2013), and 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

Table A.2: Vaccine Strategies

Type	Description	# Candidates
RNA (genetic)	Consist of messenger RNA molecules which code for parts of the target pathogen that are recognised by our immune system ('antigens'). Inside our body's cells, the RNA molecules are converted into antigens, which are then detected by our immune cells.	33
DNA (genetic)	Consist of DNA molecules which are converted into antigens by our body's cells (via RNA as an intermediate step). As with RNA vaccines, the antigens are subsequently detected by our immune cells.	21
Viral Vector	Consist of harmless viruses that have been modified to contain antigens from the target pathogen. The modified viruses act as delivery systems that display antigens to our immune cells. Replicating make extra copies of themselves in our body's cells. Non-replicating do not.	56
Protein	Consist of key antigens from the target pathogen that are recognised by our immune system.	78
Inactivated	Consist of inactivated versions of the target pathogen. These are detected by our immune cells but cannot cause illness.	16
Attenuated	Consist of living but non-virulent versions of the target pathogen. These are still capable of infecting our body's cells and inducing an immune response, but have been modified to reduce the risk of severe illness.	4

Note: Table describes the number of vaccine candidates in each strategy. 51 candidates have other, virus-like particle or unknown strategies. Data from the London School of Hygiene & Tropical Medicine's COVID-19 Tracker. Data as of November 2, 2020.

from the pharmaceutical and financial press during 2020 as detailed in the appendix above.

Table A.6 summarizes the distribution of time spent in each state in our simulation. Following Wong et al. (2018), we adopt each candidate's most advanced state. We track days spent in each state until the next state starts, only among candidates that have successfully transitioned to the next state. The realized outcomes for durations are reasonably consistent with our choices of parameters, in particular for Phase I and Phase II. And the standard deviations of durations are less than the mean is consistent with the Gaussian copula assumption of positively correlated outcomes.

We then augment π_s^{base} with 233 news articles from FactSet StreetAccount, split into positive and negative news types. Table A.7 lists the news types along with their changes in probabilities.

Table A.3: Number of Articles by News Type

News Type	Number of Articles
Release positive data	79
Announce next state	45
Positive regulatory action	30
Positive preclinical progress	22
Announce dosage start	21
Positive enrollment	17
State ahead of schedule	7
State resumed	5
State paused	4
State behind schedule	1
Negative regulatory action	1
Negative enrollment	1
Total	233

Note: Table shows the count of news articles by news type.

Table A.4: Number of Articles by Top 10 Candidates

Candidate	Number of Articles
Moderna	37
BioNTech / Fosun Pharma / Pfizer	25
Oxford / AstraZeneca	23
Johnson & Johnson / Beth Israel Deaconess Medical Center	21
Inovio Pharmaceuticals	18
Novavax	14
Arcturus / Duke	10
Vaxart	9
Medicago / GSK / Dynavax	8
Takis / Applied DNA / Evvivax	8

Note: Table the number of news articles for the top ten candidates by article count.

Table A.5: State Durations and Probabilities of Success

State	τ_s (years)	π_s^{base} (%)
Preclinical	0.6	5
Phase I	0.2	70
Phase II	0.2	44
Phase III	0.4	69
Application	0.1	88
Approval	0.5	95

Table A.6: Vaccine States

	Days in State				
	Min	Max	Mean	Median	SD
Preclinical	1.0	233.0	94.6	90.5	59.2
Phase I	17.0	103.0	51.9	27.0	39.8
Phase II	6.0	152.0	86.8	89.0	54.5

Note: Table shows statistics on days spent in each state before transitioning, among candidates that have successfully transitioned to the next state. Following Wong et al. (2018), we adopt each candidate's most advanced state. Data are from the LSHTM and are as of November 2, 2020.

Table A.7: News and Changes in Probabilities

Positive		Negative	
News type	$\Delta\pi$ (%)	News type	$\Delta\pi$ (%)
Announce next state	+5	Pause in state	-25
State ahead of schedule	+2	State behind schedule	-15
Release positive data	+5	Release negative data	-60
Positive regulatory action	+3	Negative regulatory action	-50
Positive preclinical progress	+1	Negative preclinical progress	-2
Positive enrollment	+1	Negative enrollment	-5
Dose starts	+1		
State resumes after pause	+5		

Note: Table shows the positive and negative news types, along with their changes in probabilities.

C Extensions

This appendix considers two extensions of the model: one with endogenous investment in a disaster mitigation technology, and a second with endogenous labor supply that increases exposure to the disaster shocks.

C.1 Endogenous Vaccine Development

The model in Section 2 includes no actual vaccine development technology. The reason for this is simply that parameterizing and calibrating a bio-pharmaceutical R&D production function is beyond the scope of our study. However, there could be a concern that we overstate the value of ending the pandemic by not giving the economy a real option to address it. We now show a tractable way to do this, and we explain why our results are consistent with this extension. In a nutshell, optimal research effort will impose a constraint on the parameters that does not affect our empirical identification of the pandemic duration and severity.

Suppose that, when a pandemic arrives, the representative agent has the ability to choose an expenditure rate ι that increases the arrival rate of vaccine progress. The most parsimonious specification would just be linear:

$$\lambda(\iota) = L_0 + L_1\ \iota.$$

(The discussion will treat the 2-state model. Generalization to S-states is straightforward.) Given the rate, the dynamics of wealth, dq/q, picks up a new term $-\iota\ dt$ for the duration of the pandemic. Without loss of generality, we can assert that whatever ι level the agent chooses in the first pandemic is also optimal for all subsequent pandemics. For notational simplicity below, define the adjusted drift during the pandemic as

$$\mu_S(\iota) = \mu(1) - \iota$$

where $\mu(1)$ is the benchmark growth rate without research effort. While this formulation is too sparse to address issues of public versus private returns to research expenditure, it does allow us to formulate and solve a model in which vaccine progress (and exiting from pandemic) is an endogenous outcome.

Notice that, given a choice of ι, the economy behaves exactly as in our reduced-form case. Hence the solution for optimal consumption and the value function are unchanged. In particular, we can write the value function within the pandemic as $H(1;\lambda(\iota), \mu_S(\iota))$. To choose the optimal policy, ι^*, the agent simply maximizes this function. This verifies that the optimal rate is constant within a state and does not vary across pandemics. Of course, a necessary condition for an interior optimum is

$$\frac{\partial H}{\partial \mu_S} = L_1 \frac{\partial H}{\partial \lambda}.$$

Optimality of the research effort does constrain admissible pairs of λ and μ_S via this relation.

Clearly in an economy with powerful research technology (where L_1 is a large number), agents can make the pandemic very brief (in expectation) at low cost. Hence, the endogenous value of λ would be high, and agents would pay less to return to the non-pandemic state than in an economy with inferior vaccine technology. However, recall that our benchmark calibration above already conditioned on different values of λ. We showed that the welfare gain depended strongly on the remaining expected duration of the pandemic, which could be inferred in the data from our estimation of the expected time to deployment of a vaccine during 2020.

Now, taking λ as fixed at an observed value $\hat{\lambda}$ say, consider the ratio

$$f(\hat{\mu}_S) = \left.\frac{\partial H}{\partial \mu_S}\right|_{\hat{\lambda},\hat{\mu}_S} \Big/ \left.\frac{\partial H}{\partial \lambda}\right|_{\hat{\lambda},\hat{\mu}_S}.$$

Given any value of the technology parameter L_1, the first order condition above requires us to use the value of $\hat{\mu}_S$ satisfying $f(\hat{\mu}_S) = L_1$. Assuming a solution exists, this is the full economic content of endogenizing vaccine investment in this setting.

Would imposing such a restriction on μ_S affect our estimated welfare results? To see why it will not, recall from Section 2.6 that the stock market response to news about the state (i.e., vaccine progress) effectively identifies the welfare gradient directly. The precise choice of individual model parameters is, to a first approximation, irrelevent, conditional on matching this moment. Requiring that μ_S satisfies the above first order condition would take away one degree of freedom in the calibration. But choosing the remaining disaster parameters (e.g. χ) to yield the same stock market response would restore the same empirical conclusions.

To be clear, the conclusion is not that including a vaccine development technology does not affect the welfare costs of the pandemic. Rather, we are pointing out that our empirical work has already pinned down the key inputs to that value. Taking those quantities at face value, adding assumptions about the development technology and imposing the restriction of optimal investment do not perturb the calculation.

C.2 Endogenous Pandemic Severity and Labor Externalities

In this subsection, we develop a version of the model in which the pandemic parameters for the wealth process are endogenized through the choice of labor supply. Doing so will allow us to examine how much the value of curtailing the pandemic is influenced by the extent to which individual choices deviate from the socially optimal policies.

In this version of the model, wealth accumulates according the stochastic process

$$dq = \ell^{\alpha} q\mu dt - Cdt + \sigma \ell^{\alpha/2} q dB_t \tag{A.51}$$

in the non-pandemic state, and

$$dq = \ell^{\alpha} q\mu dt - Cdt + \sigma \ell^{\alpha/2} q dB_t - [\ell\varepsilon + k + KL] q dJ_t. \tag{A.52}$$

in the pandemic state. As before, C is the endogenous consumption rate, and now ℓ is the household's labor supply, and $\alpha \in (0,1)$ is the elasticity of expected output with respect to labor. The results below all go through with constant returns to scale in the drift term. Crucially, both individual and aggregate labor are assumed to affect the agent's exposure to the health shock via the jump size. Let

$$\chi(\ell, L) \equiv [\ell\varepsilon + k + KL], \tag{A.53}$$

where ε is exposure to the pandemic via private labor, k is exposure to the pandemic unrelated to labor, L is aggregate labor supply, and K is exposure via aggregate labor. These parameters can capture losses of wealth due to health-induced disruptions to work, the need to work from home with attendant productivity impact and loss of human capital, deadweight losses from bankruptcy, and frictions from labor reallocation. We will assume parametric restrictions on ε, k and K to be small enough that $(1-\chi) \in (0,1)$. The agent takes the aggregate supply of labor L as given in her optimization problem.

Agents' preferences are as in the text. We assume no disutility to labor supply and no frictions in adjusting ℓ. We assume $\ell \in [0, \overline{\ell}]$, where the upper bound $\overline{\ell}$ is the agents' total available work capacity. (In the numerical work we normalize $\overline{\ell} = 1$.)

The agent's problem is now to choose in each state s optimal consumption $C(s, L^*(s))$ and labor $\ell(s, L^*(s))$ that maximizes the objective function. We impose that agents have rational expectations about $L^*(s)$, the aggregate labor in equilibrium. In other words, individual agents' decisions in the aggregate should lead to a wealth (consumption) dynamic that is confirmed in equilibrium. This implies the following for wealth dynamics in the pandemic regime:

$$dq(s) = [\ell(s, L^*(s))]^\alpha q\mu dt - C(s, L^*(s))dt + \sigma[\ell(s, L^*(s))]^{\alpha/2} qdB - \chi(\ell(s, L^*(s)), L^*(s))qdJ_t \tag{A.54}$$

Since $L^*(s)$ is a constant for each s, as the agent has rational expectations about $L^*(s)$, the above dynamics are identical to those assumed by the agent. Substituting for the equilibrium fixed point that $L^*(s) = \ell(s, L^*(s))$, we can then obtain the rational expectations equilibrium outcomes.

Proposition 1. *Equilibrium labor in the non-pandemic state is given by*

$$L(0) = L(S) = \bar{\ell} \tag{A.55}$$

Equilibrium labor in pandemic states $L^*(s)$ $\forall s \in \{1, \ldots, S-1\}$ solves[3]

$$\chi(L(s), L(s)) = k + (\varepsilon + K)L(s) = \left[1 - (L(s))^{\frac{1-\alpha}{\gamma}}\nu\right] \tag{A.56}$$

where

$$\nu \equiv \left[\frac{\alpha\left(\mu - \frac{1}{2}\gamma\sigma^2\right)}{\zeta\varepsilon}\right]^{-\frac{1}{\gamma}}. \tag{A.57}$$

Proof. The HJB equation for each state $s \in \{1, \ldots, S-1\}$ is now

$$\begin{aligned} 0 = \max_{C,\ell} \Big[& f(C, \mathbb{J}(s)) - \rho\mathbb{J}(s) + \mathbb{J}_q(s)(\ell^\alpha q\mu - C) + \frac{1}{2}\mathbb{J}_{qq}(s)\ell^\alpha q^2\sigma^2 + \zeta\left[\mathbb{J}(s)\left(q(1-\chi)\right) - \mathbb{J}(s)(q)\right] \\ & + \lambda_u(s)\left[\mathbb{J}(s+1)(q) - \mathbb{J}(s)(q)\right] + \lambda_d(s)\left[\mathbb{J}(s-1)(q) - \mathbb{J}(s)(q)\right]\Big] \end{aligned} \tag{A.58}$$

[3] It can be shown that given $\alpha \in (0, 1)$, the second order condition for a maximum is satisfied whenever $\mu - \frac{1}{2}\gamma\sigma^2 > 0$, which also implies $\nu > 0$.

Using the conjecture for the objective function in the text for $\mathbb{J}(s)$, calculating the derivatives with respect to q, $\mathbb{J}_q(s) = H(s)q^{-\gamma}$ and $\mathbb{J}_{qq}(s) = -\gamma H(s)q^{-\gamma-1}$, and differentiating with respect to labor ℓ, we obtain the first-order condition as

$$\mathbb{J}_q(q)\alpha\ell^{\alpha-1}\mu q + \frac{1}{2}\mathbb{J}_{qq}(q)\alpha\ell^{\alpha-1}\sigma^2 q^2 - \mathbb{J}_q\left(q(1-\chi)\right)\zeta\varepsilon q = 0 \tag{A.59}$$

where we have suppressed state s in the notation. This in turn simplifies to

$$\left[\frac{\alpha\left(\mu - \frac{1}{2}\gamma\sigma^2\right)}{\zeta\varepsilon}\right]\ell^{\alpha-1} - [1-\chi]^{-\gamma} = 0 \tag{A.60}$$

where $\chi(\ell, L) = k + \varepsilon\ell + KL$. In rational expectations equilibrium $L(s) = \ell(s)$, which gives us that optimal labor in pandemic state $L^\star(s)\ \forall s \in \{1, \ldots, S-1\}$ satisfies (A.56):

$$\chi\left(L(s), L(s)\right) = k + (\varepsilon + K)L(s) = \left[1 - \left(L(s)\right)^{\frac{1-\alpha}{\gamma}}\nu\right] \tag{A.61}$$

where

$$\nu \equiv \left[\frac{\alpha\left(\mu - \frac{1}{2}\gamma\sigma^2\right)}{\zeta\varepsilon}\right]^{-1/\gamma}. \tag{A.62}$$

The second-order condition with respect to ℓ is satisfied (see footnote 7 above) whenever $\left(\mu - \frac{1}{2}\gamma\sigma^2\right) > 0$. For the non-pandemic state $s = 0$ or $s = S$, the third term in first-order condition (A.59) is absent; therefore, we obtain that labor is at the highest possible level $L(0) = L(S) = \overline{\ell}$, whenever $\alpha\left(\mu - \frac{1}{2}\gamma\sigma^2\right) > 0$. □

In the non-pandemic state, the agent faces no cost to supplying labor and exerts effort fully. However, in the pandemic states, the agent increases exposure to health risk by supplying labor, which creates a tradeoff between augmenting the capital stock and reducing the loss of capital that arises from health shocks. A key property of the model is that the agent contracts labor relative to the non-pandemic state.

Note the externality in our set up via the KL term in the size of the Poisson shock (where L is aggregate labor) that is not internalized by each agent. A central planner would factor this in the socially efficient choice of labor. This is tantamount to replacing ε by $(\varepsilon + K)$ in ν above to obtain ν^{CP}:

$$\nu^{CP} \equiv \left[\frac{\alpha \left(\mu - \frac{1}{2}\gamma\sigma^2 \right)}{\zeta(\varepsilon + K)} \right]^{-\frac{1}{\gamma}} \tag{A.63}$$

Socially efficient labor choice $L^{CP}(s)$ in the pandemic states is then given by

$$\chi(L(s), L(s)) = k + (\varepsilon + K)L(s) = \left[1 - (L(s))^{\frac{1-\alpha}{\gamma}} \nu^{CP} \right] \tag{A.64}$$

It is then straightforward to show that $\nu^{CP} > \nu$ for $K > 0$ and $\gamma > 0$, and hence $L^{CP}(s) < L(s)$, i.e., the socially efficient choice of labor in pandemic states is smaller than the privately optimal one.

Given the optimal labor and consumption policies, the model solutions in Proposition 2 can be directly applied. As before, the pandemic parameters only enter the system of equations via the constants g_0 and g_1, which we can write compactly as

$$g(x,y) \equiv \frac{(1-\gamma)\rho}{(1-\psi^{-1})} - x^{\alpha}(1-\gamma)\left(\mu - \frac{1}{2}\gamma\sigma^2\right) - y\left([1-\chi(x,x)]^{1-\gamma} - 1\right) \tag{A.65}$$

with $g_0 = g(\bar{\ell}, 0)$ and $g_1 = g(\ell(s), \zeta)$.

To quantitatively evaluate the model's implications, we require that the parameters are such that the endogenous severity of the pandemic is in line with our empirical estimates. A useful summary statistic for the severity is the reduction in the expected growth of wealth (computed under the risk-neutral measure) which we label Δm_Q. For reference, under the basline calibration in the text, the value of this quantity is approximately 0.06. We report below the implied Δm_Q for a range of values of K and ε in Table A.8.[4] We also report the optimal labor supply in the pandemic state, $\ell^{\star}$. Some empirical evidence suggests labor contraction ≈20% in April 2020 (see, e.g., Cajner et al. (2020)) corresponding to $\ell^{\star} \approx 0.80$. The table identifies parameter regions (e.g., the upper left of the tables) that can match both restrictions.

Table A.9 shows the effect on V of the labor market externality for the same range of parameter values. The left panel provides a direct measure of the scale of the externality via the

[4]The exercise fixes α and k. These parameters have less direct impact on the degree of labor externality.

Table A.8: Endogenous Pandemic Parameters via Labor

	$\ell^\star$				Δm_Q		
	$K \rightarrow$				$K \rightarrow$		
$\varepsilon \downarrow$	0.8471	0.8204	0.7959	$\varepsilon \downarrow$	0.0517	0.0577	0.0636
	0.7885	0.7651	0.7435		0.0509	0.0565	0.0619
	0.7357	0.7151	0.6960		0.0503	0.0555	0.0605
	0.6880	0.6698	0.6529		0.0497	0.0546	0.0593

Note: Table shows the implied values of equilibrium labor, $\ell^\star$, and decline in expected growth rate of q in pandemic states, Δm_Q, by fixing the elasticity of expected output with respect to labor $\alpha = 0.5$, exposure to the pandemic unrelated to labor $k = 0.006$, intensity of switching to the pandemic state $\eta = 0.04$, and intensity of switching to the non-pandemic state $\lambda = 0.5$. Each panel varies the exposure to the pandemic via private labor, ε, and via aggregate labor, K. ε increases down the rows and takes the values 0.023, 0.024, 0.025 and 0.026, while K increases left-to-right across columns and takes the values 0.018, 0.024 and 0.030.

ratio of the central planner's solution for optimal labor in the pandemic to that actually chosen by agents. With parameters in the region identified above, the socially optimal lockdown is quite severe with labor restricted to 30%-40% of the privately optimal amount.[5] The right panel shows that, in this region, the welfare gain is 12%-19% lower under the central planner's solution.

In addition to the finding that ending the pandemic is less valuable under a central planner, comparing variation across the two panels reveals the pattern that a stronger externality (as measured by lower values of $\ell^\star_{cp}/\ell^\star$) are associated with decreasing relative welfare gains under the central planner. The extra degree of lockdown that the planner would impose decreases the expected welfare cost. We acknowledge that if the arrival of the pandemic were to result in social costs that are outside the capital stock dynamics for the agent, then the planner might value the ending the pandemic more than the representative agent.

[5]While our model does not feature SIR dynamics, models with SIR dynamics and labor externalities generally see more severe lockdown policies under a central planner (see Abel and Panageas (2021)).

Table A.9: Externality and Value of a Cure

	$\ell^{\star}_{cp}/\ell^{\star}$				V_{cp}/V		
		$K \rightarrow$				$K \rightarrow$	
$\varepsilon \downarrow$	0.3762	0.2994	0.2450	$\varepsilon \downarrow$	0.8661	0.8167	0.7725
	0.3860	0.3083	0.2530		0.8777	0.8309	0.7884
	0.3957	0.3170	0.2609		0.8882	0.8438	0.8031
	0.4049	0.3256	0.2686		0.8975	0.8555	0.8165

Note: The left panel provides a direct measure of the scale of the externality via the ratio of the central planner's solution for optimal labor in the pandemic to that actually chosen by agents. The right panel shows the ratio of the value of a cure as determined by the central planner to that chosen by agents. Both fix the elasticity of expected output with respect to labor $\alpha = 0.5$, exposure to the pandemic unrelated to labor $k = 0.006$, intensity of switching to the pandemic (non-pandemic) state $\eta = 0.4$ ($\lambda = 0.5$). Exposure to the pandemic via private labor, ε, increases down to rows and takes values $\{0.023, 0.024, 0.025, 0.026\}$. Exposure via aggregate labor, K, increases left-to-right and takes values $\{0.018, 0.024, 0.030\}$.

D Interpreting the COVID-19 stock market experience

For tractability, our model omits many factors that played important roles in 2020. We have not included fiscal or monetary policy, for example. (Recall, though, that our empirical work excluded large market-moving events attributable to non-vaccine news as classified in Baker et al. (2020)). We now describe in more detail its implications for the response of the stock market and of consumption to a pandemic. Comparing these implications to the actual experience of 2020, we argue that a coherent interpretation is possible when viewed through the lens of the model. Hence, while there are inevitably descriptive limitations, these need not detract from the paper's objectives.

Consider first the **stock market**. From December 31, 2019 to March 23, 2020 the CRSP value-weighted experienced a return of approximately -36%, continuously compounded. The market then rebounded fully by early autumn. The cumulative return for the year was approximately +5% at end of October (where our clinical trial sample stops).

The calibration in the paper implies a somewhat smaller drop, of -25% for the output claim at the onset of a pandemic, based on our initial VPI forecast of an expected duration of four years. Thus, relative to this data point, our estimation of the potential damage of the pandemic is conservative compared to the market's assessment.

Subsequently, the model implies a partial market recovery due to the observed success of vaccine trials. From March 23 through October 30, our forecast of the pandemic's duration dropped by 2.5 years, of which 0.6 years was expected. The implied market response to this progress – calibrated to match the response estimated in our empirical work – is approximately 16%.

Thus, without conditioning on any other shocks, the basic dynamics of our model can account for approximately 70% of the observed market decline and about one-third of the recovery through our sample period. To shed light on factors the model may be missing, it is useful to decompose the initial market decline into components due to cash flow news, real interest rate news, changes in the equity risk premium. One way to identify these return components, r^{CF}, r^{RF},and r^{RP} within the model is as follows.

1. Assume that, on a switch to a pandemic, the parameters of the process dq change as described in Section 2. Solve the asset pricing system given in Proposition 3, but without any risk-adjustments (i.e., under the physical measure) and fixing the riskless rate in the pandemic, $r(1)$, to be its non-pandemic value, $r(0)$, as computed under the original benchmark calibration.

    - Call the resulting price-dividend ratio $\tilde{p}(s)$.
    - Define the percent change in this on entering the pandemic $r^{CF} = \log(\tilde{p}(1)/\tilde{p}(0))$ to be the return due to cash flow news.

2. Solve the same pricing system, still under the physical measure, but now allowing $r(1)$ to be its pandemic value in the calibrated model.

    - Call the new price-dividend ratio $\hat{p}$.
    - Call $r^{RF} = \log(\hat{p}(1)/\hat{p}(0)) - r^{CF}$ the return due to real interest rate news.

3. As in the text, let $p(s)$ denote the price-dividend ratio under the full model.

    - Call the residual $r^{RP} = \log(p(1)/p(0)) - r^{RF} - r^{CF}$ the return due to risk premium news.

With the baseline calibration the results are:

$$r^{CF} = -0.328, \qquad r^{RF} = 0.284, \qquad r^{RP} = -0.206.$$

Recently, Knox and Vissing-Jorgensen (2022) report empirical estimates of the same decomposition during 2020 using information in options markets, inflation swaps, and S&P 500 dividend futures. Consistent with their work, we find substantial and nearly offsetting positive and negative components of discount rate news, that is $r^{RF} + r^{RP}$ is small.

The risk premium component in our model is due to the threat of Poisson shocks. As is well-known from the long-run risk literature, investors with Epstein-Zin preferences and $\gamma > 1/\psi$ are averse to uncertainty and growth-rate risk. Also, in the model, real rates decline on entering the pandemic with the riskless rate turning mildly negative, consistent with the data.

Our model attributes large effects to cash flow news. In fact, expected cash flows did decline steeply in early 2020: December 2020 S&P 500 dividend futures declined *more* in percentage terms than did the overall stock market,[6] However, as Knox and Vissing-Jorgensen (2022) show, near-term dividends makes up a small component of market value. So declines in the discounted sum (e.g., for 10 years ahead) cannot account for a large component of the market drop. As a result, and combined with the small net discount rate effects, the authors attributing most of the market decline to residual unidentified long-term effects.

In our model, there are both short-term cash-flow effects from lower output, as well long-term effects from the loss of part of the capital stock. Evidently the market was anticipating less short-run impact, but greater long-term impact, perhaps due to scarring type effects that are absent from the model. Recall that, while we assume permanent effects of the pandemic on the level, q, we assume purely transitory effects on the growth rate of dq once the pandemic ends. The market may have not been so sure.

Turning to the implication for our conclusions, it is clear that adding negative long run effects in order to match the observed market decline (while holding other return components fixed) will imply larger welfare gains to mitigating the pandemic. It is important to recognize that, while the model's return decomposition can be altered (e.g., with different preference

[6] The December contract, quoted in units of the S&P 500 index, dropped from 62.5 at the end of 2019 to 39 on 3/23/20, a decline of 47%. Using this contract as a denominator, the price-dividend ratio on the market actually went up over this period.

parameters), this will not necessarily lead to large welfare effects if we still impose that the calibration matches the magnitude of the market response to vaccine progress. As described in the text, the latter condition effectively pins down the severity of the pandemic.

Next consider the path of **consumption**. An extremely prominent feature of the 2020 experience was the rapid plunge in consumption in March and April followed by a complete rebound by early 2021. In the context of the model, consumption is driven primarily by the wealth process, dq. (The effect of changes in the marginal propensity to consume is second order.) A large decline in observed consumption is consistent with the occurrence of one or more down jumps in q in the early part of the year. These Poisson events are the way in which the pandemic is realized within the model. In fact, the early occurrence of such a shock – which was not considered above – can also bring the model's implied stock market decline directly into line with the observed fall.

However, as just discussed, the model has no mechanism for the reversal of these shocks after the pandemic ends, still less while the pandemic remains in progress. How, then, can the model explain the consumption rebound?

In our view, the most coherent interpretation of the consumption experience is that a large component of the recovery must be regarded as having been unexpected. Indeed, from the analysis above, it would be seemingly impossible to build a model in which a strong rebound in consumption is expected *ex ante* within a pandemic <u>and</u> in which the stock market drops nearly 40% due to the arrival of the pandemic. Moreover, the strong rebound in stock prices after March is also consistent with the interpretation of substantial unexpected good news about fundamentals, further helping to reconcile the model's implications of a rally only partially explained by vaccine progress. Moreover, evidence in Hong et al. (2021) and Gormsen and Koijen (2020), respectively, supports strong positive revision in corporate cash flows during the pandemic, by examining expectations of stocks' earnings and implied dividend yields.

In the context of the model, unexpected consumption changes are described by the Gaussian component of wealth shocks, which can be viewed as encompassing mechanisms like (unanticipated) policy interventions that are outside the model. Invoking large positive shocks is not implausible if the scale of these shocks, governed by $\sigma(s), s > 1$, is large. Our calibration too $\sigma(s) = 0.075$ (or 3.75% per quarter) for the pandemic states. With this value, the 8% increase in consumption in the third quarter of 2020 is approximately a two standard deviation event (depending on the assumed conditional mean), unlikely but not impossible.

To summarize, fully accounting for the behavior of the financial markets and the real economy during 2020 is beyond the scope of our baseline stylized model. Nor is it the main objective of the paper. Nevertheless, the primary contours of the data can be reasonably described as an outcome within the model, given certain realizations of the stochastic shocks, as described above.

References

Andrew Abel and Stavros Panageas. Social distancing, vaccination and the paradoxical optimality of an endemic equilibrium. Technical report, Working Paper, 2021.

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 2020.

Tomaz Cajner, Leland Dod Crane, Ryan Decker, Adrian Hamins-Puertolas, and Christopher J Kurz. Tracking labor market developments during the covid-19 pandemic: A preliminary assessment. Technical report, 2020.

Niels Joachim Gormsen and Ralph S J Koijen. Coronavirus: Impact on Stock Prices and Growth Expectations. *The Review of Asset Pricing Studies*, 10(4):574–597, 09 2020.

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 2021.

Benjamin Knox and Annette Vissing-Jorgensen. A stock return decomposition using observables. Working paper, FEDS Working Paper, March 2022.

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.

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.