

DISCUSSION PAPER SERIES

DP15558

(v. 3)

The Value of a Cure: An Asset Pricing Perspective

Viral Acharya, Tim Johnson, Suresh M Sundaresan
and Steven Zheng

FINANCIAL ECONOMICS

MACROECONOMICS AND GROWTH

PUBLIC ECONOMICS

CEPR

The Value of a Cure: An Asset Pricing Perspective

Viral Acharya, Tim Johnson, Suresh M Sundaresan and Steven Zheng

Discussion Paper DP15558
First Published 14 December 2020
This Revision 19 March 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programmes:

- Financial Economics
- Macroeconomics and Growth
- Public Economics

Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Viral Acharya, Tim Johnson, Suresh M Sundaresan and Steven Zheng

The Value of a Cure: An Asset Pricing Perspective

Abstract

We provide an estimate of the value of a cure using the joint behavior of stock prices and a vaccine progress indicator during the ongoing COVID-19 pandemic. Our indicator is based on the chronology of stage-by-stage progress of individual vaccines and related news. We construct a general equilibrium regime-switching model of repeated pandemics and stages of vaccine progress wherein the representative agent withdraws labor and alters consumption endogenously to mitigate health risk. The value of a cure in the resulting asset-pricing framework is intimately linked to the relative labor supply across states. The observed stock market response to vaccine progress serves to identify this quantity, allowing us to use the model to estimate the economy-wide welfare gain that would be attributable to a cure. In our estimation, and with standard preference parameters, the value of the ability to end the pandemic is worth 5-15% of total wealth. This value rises substantially when there is uncertainty about the frequency and duration of pandemics. Agents place almost as much value on the ability to resolve the uncertainty as they do on the value of the cure itself. This effect is stronger – not weaker – when agents have a preference for later resolution of uncertainty. The policy implication is that understanding the fundamental biological and social determinants of future pandemics may be as important as resolving the immediate crisis.

JEL Classification: G12, D5, I1, Q54

Keywords: Pandemic, Vaccine, COVID-19, Rare Disasters, Regime-Switching, parameter uncertainty

Viral Acharya - vacharya@stern.nyu.edu
Stern School of Business, NYU and CEPR

Tim Johnson - tcj@illinois.edu
University of Illinois at Urbana Champaign

Suresh M Sundaresan - ms122@columbia.edu
Columbia University

Steven Zheng - steven.zheng@stern.nyu.edu
New York University Stern School of Business

Acknowledgements

We thank Dick Berner, Rob Engle, Matt Richardson, Venky Venkateswaran, and Olivier Wang for their comments and suggestions. We also received valuable comments from participants at NYU Stern Finance Department seminar, Advisory Board Meeting of NYU Stern Volatility and Risk Institute, and UIUC Gies Finance Department seminar. We are grateful to the Vaccine Centre at the London School of Hygiene & Tropical Medicine for sharing data.

The Value of a Cure: An Asset Pricing Perspective*

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

March 19, 2021

Abstract

We provide an estimate of the value of a cure using the joint behavior of stock prices and a vaccine progress indicator during the ongoing COVID-19 pandemic. Our indicator is based on the chronology of stage-by-stage progress of individual vaccines and related news. We construct a general equilibrium regime-switching model of repeated pandemics and stages of vaccine progress wherein the representative agent withdraws labor and alters consumption endogenously to mitigate health risk. The value of a cure in the resulting asset-pricing framework is intimately linked to the relative labor supply across states. The observed stock market response to vaccine progress serves to identify this quantity, allowing us to use the model to estimate the economy-wide welfare gain that would be attributable to a cure. In our estimation, and with standard preference parameters, the value of the ability to end the pandemic is worth 5-15% of total wealth. This value rises substantially when there is uncertainty about the frequency and duration of pandemics. Agents place almost as much value on the ability to resolve the uncertainty as they do on the value of the cure itself. This effect is stronger – not weaker – when agents have a preference for later resolution of uncertainty. The policy implication is that understanding the fundamental biological and social determinants of future pandemics may be as important as resolving the immediate crisis.

JEL Codes: G12, D5, I1, Q54

Keywords : pandemic, vaccine, COVID-19, rare disasters, regime-switching, parameter uncertainty

*First draft: November 15, 2020. We thank Dick Berner, Rob Engle, Matt Richardson, Venky Venkateswaran, and Olivier Wang for their comments and suggestions. We also received valuable comments from participants at NYU Stern Finance Department seminar, NYU Stern Quantitative Finance and Econometrics (QFE) seminar, Advisory Board Meeting of NYU Stern Volatility and Risk Institute, and UIUC Gies Finance Department seminar. 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.

[†]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.

[¶]New York University, Stern School of Business: steven.zheng@stern.nyu.edu

1 Introduction

Quantifying the scale of the economic damage caused by a pandemic is a crucial step in assessing policy responses along social, medical, fiscal, and monetary dimensions. This paper builds on the hypothesis that stock markets may contain valuable information for gauging the value of *ending* the pandemic. Stock markets, which corrected by as much as 40-50% at the outbreak of the coronavirus pandemic in February-March 2020, have rebounded robustly within six months. While there are many explanations proposed for the seeming disconnect between the real economy ravaged by the pandemic and the buoyant stock market, one candidate on the table relates to the progress in development of vaccines¹ to end the pandemic. On the one hand, only the arrival (and delivery) of an efficacious vaccine is considered as a definitive event that will end the pandemic and result in robust economic recovery.² On the other hand, stock prices – by reflecting forward-looking expectations – should reflect the economic value of credible progress in the development of vaccines; this value arises from the ability of vaccines to end the pandemic and is naturally related to the scale of the economic damage caused by the pandemic.

The relationship between stock prices and vaccine development is well-illustrated by the following examples. On May 18 and July 14, 2020, *Moderna*, one of the vaccine developing companies, announced good news relating to the progress in its Phase I clinical trials and moving to the next stage of trials. Similarly, on November 9, 2020, *Pfizer* and *BioNTech* announced positive news regarding their Phase III clinical trials. In response to these news, the U.S. stocks gained over \$1 trillion in cumulative market capitalization over these three days, with several pandemic-exposed sectors such as airlines, cruise ships, and hotels experiencing 10-20% appreciations on each day. These moves were both economically large and indicative of time to deployment of a vaccine being an important factor driving variation in stock market prices.³

We build upon these observations and offer an asset-pricing perspective to estimate the value of a cure, i.e., the amount of wealth that a representative agent would be willing to pay for obtaining a vaccine that puts an end to the pandemic. While there are now several estimations in

¹We use “cure” and “vaccine” interchangeably to denote something that brings the pandemic to an end, despite being medically very different.

²See [Lauren Fedor and James Politi, Financial Times, May 18, 2020](#) in the Appendix.

³See (1) [Matt Levine, Money Stuff, May 19, 2020](#), (2) [Matt Levine, Money Stuff, July 16, 2020](#), (3) [John Authers, Bloomberg Opinion, November 10, 2020](#), and (4) [Laurence Fletcher and Robin Wigglesworth, Financial Times, November 14, 2020](#) in the Appendix.

the literature of how costly the pandemic is to the economy, our approach is different in that it uses stock market data to calculate the *ex-ante* value of a cure. Our approach is directly analogous to the seminal work of Lucas (1987) in assessing the welfare costs associated with business cycle risk. Just as that paper provides a framework for assessing the consequences of policy responses to mitigate output volatility, our work speaks to the cost-benefit analysis of potential public sector investment in alleviating the threat of current and future pandemics. Going beyond our baseline estimate for the value of a cure, we can also shed light on which of the primitive elements of the economy are driving that value.

Our analysis proceeds as follows.

First, we document empirically the joint behavior of stock returns (for market portfolio and cross-section of industries) and expected time to deployment of a vaccine. To this end, we construct a novel "vaccine progress indicator." Our indicator is based on the chronology of stage-by-stage progress of individual vaccines (obtained from the Vaccine Centre at the London School of Hygiene & Tropical Medicine) and related news (obtained from FactSet). Using data on vaccine development for past epidemics and surveys during the COVID-19 pandemic, we calibrate the probabilities of transition across different stages of vaccine development and use news to "tap" these probabilities up or down. We then simulate over 200 vaccine "trials" corresponding to the vaccines being developed, factoring in a correlation structure between trials based on relevant characteristics such as their approach ("platform"), belonging to a common company, etc. The result of this exercise is a vaccine progress indicator using all available information at a given point of time expressed in terms of expected time to deployment of a vaccine.⁴ The evolution of our indicator is shown in Figure 1.

We then relate stock market returns to changes in the expected time to deployment of a vaccine by regressing the returns on changes in our vaccine progress indicator, controlling for lagged returns as well as large moves attributable to release of other macroeconomic news. Allowing for some lead-lag structure in the relationship, e.g., due to leakage of news or dating noise in our news data, we estimate that a reduction in the expected time to deployment of a vaccine by a year results in an increase in the stock market return as a whole by between 4 to 8% on a daily basis.

⁴An 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.

The joint relationship exhibits the anticipated cross-sectional properties, with the co-movement between returns and changes in the vaccine progress indicator being stronger for sectors most affected by the COVID-19 pandemic (see Figure 4).

Second, we build a general equilibrium regime-switching model of pandemics with asset pricing implications to translate this empirical co-movement of stock returns and vaccine progress indicator into the value of a cure. We consider an economy with a representative agent who has stochastic differential utility (Epstein-Zin preferences) with endogenous labor and consumption choices. The state of the economy can be “normal,” i.e., without a pandemic, or in a pandemic; within the pandemic, there are several regimes mapping into the stages of vaccine development. The economy transitions across these states based on a set of stationary probabilities. Once the economy switches out of a pandemic, another pandemic may occur in future. Labor augments agent’s capital stock that can be readily converted into consumption; however, labor exposes the agent to the pandemic in that within the pandemic regime, the agent can be hit by a health shock that destroys forever a part of the agent’s capital stock, and this likelihood is proportional to the labor supply.⁵ A key feature of the model is that the agent withdraws labor in the pandemic states in order to mitigate the economic exposure to a health shock. In other words, the arrival of a pandemic and the incidence of a health shock for the agent within the pandemic can be considered as “rare disasters”, the exposure to which is partly controlled by the agent.

Third, we characterize the solution to the agent’s problem of choosing labor and consumption in each state of the economy and derive the respective objective function values, which are inter-dependent but are amenable to a straightforward numerical solution of a fixed-point problem. We can then examine the pricing kernel and asset prices in this framework; in particular, we evaluate the value of a claim to future output, and study its relationship with the expected time to switching out of a pandemic state. This provides a theoretical counterpart to our empirical estimate of co-movement between stock market return and changes in vaccine progress indicator. A principal insight of our asset-pricing perspective is the following: the improvement in the welfare of the agent in switching out of a pandemic is related to the extent of contraction in labor in

⁵The permanent loss of capital stock can be due to a variety of factors such as loss of life, reduced productivity or attrition of human capital in working from home amidst closures of schools and lack of child care support, filing of bankruptcies with deadweight losses in asset value, and firing of labor with difficulty in re-matching to available jobs at a future date.

the pandemic state relative to the non-pandemic one; this same labor contraction is an important statistic (modulated by preference and pandemic parameters) that determines how sensitive are stock prices to progress towards deployment of a vaccine. The model delivers the implication that the value of moving from a pandemic state to a non-pandemic state is simply the ratio of marginal propensity to consume in the pandemic state to that in the non-pandemic state, augmented by the intertemporal elasticity of substitution. Thus, the desire to resolve uncertainty sooner is informed by the endogenous consumption choices made by the household in pandemic states.

We can therefore readily connect our empirical work to the theoretical asset-pricing perspective. With standard preference parameters employed in the literature, the value of a cure turns out to be worth 5-15% of wealth (formally, capital stock in our model). At our baseline assumptions for other parameters, this corresponds to an approximately 25% contraction of labor during the pandemic relative to the non-pandemic state. The reason why the economy would attach such a large value to the vaccine is because the pandemic causes a permanent loss of capital stock when it effects agents, which in turn is reflected in the significant precautionary contraction of labor during the pandemic. In spite of the simplicity of our model of the pandemic, we can readily examine externalities in the setup. Specifically, the representative agent can impose through its labor choice exposure for all other agents in the economy, but not internalize this spillover; we examine the difference in the value of a cure with a precautionary labor choice being made by the representative agent versus that by the central planner. Since the planner contracts the labor more and optimally reduces pandemic exposure for the economy as a whole, the planner attaches a lower value to the cure than the representative agent does.⁶

Finally, our estimate of the value of a cure depends crucially on the frequency and the expected duration of the pandemic, which raises the natural question of parameter uncertainty around these pandemic properties. Such uncertainty is natural given the rare nature of such pandemics and the evolving understanding of connections between various pandemics (SARS, H1N1, COVID-19, etc.).⁷ The final exercise we undertake is to assess the effect of imperfect information

⁶Note, however, that the planner may attach a higher value to the cure if the arrival of the pandemic were to result in social costs outside the capital stock dynamics for the agent.

⁷See, for example, "COVID-19 Is Bad. But It May Not Be the 'Big One'", Maryn McKenna, *Wired*, June 17, 2020, "Coronavirus Response Shows the World Is Not Ready for Climate-Induced Pandemics", Jennifer Zhang, *Columbia University Earth Institute*, February 24, 2020, and "The next pandemic: where is it coming from and how do we stop it?", Leslie Hook, *Financial Times*, October 29, 2020.

on the value of a cure.

We specialize our framework to just two states, non-pandemic and pandemic (without individual stages of vaccine development), but allow for uncertainty about frequency and duration of pandemics. The agent learns about these parameters as pandemics arrive and end. We can extend our asset-pricing framework also to this setting with uncertainty and learning. It turns out that the value of the cure rises sharply when there is uncertainty about the frequency and duration of pandemics. Indeed, we find that the representative agent would be willing to pay as much for resolution of this parameter uncertainty as for the cure absent such uncertainty. This effect is stronger – not weaker – when agents have a preference for later resolution of uncertainty (formally, an elasticity of intertemporal substitution, or EIS, which is lower than the inverse coefficient of relative risk aversion) as this induces a more significant contraction of labor during pandemics. An important policy implication is that understanding the fundamental biological and social determinants of future pandemics, for instance, whether pandemics are related to zoonotic diseases triggered more frequently by climate change, may be as important to mitigating their economic impact as resolving the immediate pandemic-induced crisis.

A few caveats are in order. Our model of pandemics is close to that of rare disasters in asset-pricing literature (Barro, 2006; Gabaix, 2012; Tsai and Wachter, 2015) but with endogenous exposure of the agent to disasters as well as featuring endogenous consumption, labor and asset prices. However, we do not feature SIR-style dynamics of an individual pandemic itself. Finally, our model also does not feature the impact of economic-stabilization policies such as fiscal or monetary support; however, we do control for the announcement effect of such policies in our empirical work.

The rest of the paper is organized as follows. Section 2 relates to the existing literature. Section 3 describes the construction of vaccine progress indicator and estimates of its covariance with stock market returns. Section 4 presents our general equilibrium regime-switching model of pandemics with endogenous labor and consumption decisions, and derives asset prices in this framework that help estimate the value of a cure for the pandemic. Section 5 extends (two-state version of) the model to allow for parameter uncertainty and learning to study the impact on the value of a cure from such uncertainty and the value attached to resolving it. Section 6 concludes with some

further directions for research. All proofs not in the main text are contained in the Appendix.

2 Related Literature

As noted in the introduction, our approach to quantifying the cost of the pandemic parallels that of Lucas (1987) in assessing the cost of business cycles. While Lucas (1987) finds small welfare improvements to reducing this risk, 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 have modeled climate risk using the approach that long-run risk of climate risk can manifest itself through Poisson shocks to the capital stock, which is the approach we are pursuing here in the context of pandemics. A detailed survey of this literature is provided by Tsai and Wachter (2015) for better understanding asset pricing puzzles. 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.

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. We first relate to the theoretical literature in asset pricing that is closest to our model; we then relate to the empirical literature on observed contraction in employment and consumption during the pandemic.

Hong et al. (2020b) 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. (2020a) combine the model of Hong et al. (2020b) 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 pandemics – which allows for several sub-states in a pandemic relating to vaccine progress – is an important theoretical contribution of our paper. We build upon this setup further to introduce learning when there is parameter uncertainty about pandemic parameters.⁸

For empirical work, Hong et al. (2020b) fix expected pandemic duration around one year but show in comparative statics that asset prices show considerable sensitivity to the arrival rate of the vaccine. Hong et al. (2020a) use their model to infer the arrival rate of the vaccine. In contrast, we provide a “vaccine progress indicator” in the form of an estimated time to vaccine deployment using actual data and related news on the progress of clinical trials of vaccines for COVID-19. We relate this vaccine progress indicator to stock market returns to infer labor contraction in the pandemic relative to the non-pandemic state.

Elenev et al. (2020) incorporate a pandemic state with low but disperse firm productivity that recurs with low probability for studying government intervention in corporate credit markets. While we do not model credit markets in our setup, our differentiating novel features are: construction of a vaccine progress indicator and estimation of its joint relationship with stock markets, and mapping it into a general equilibrium regime-switching model of pandemics with asset prices in order to derive an estimate of the value of a cure.

Kozlowski et al. (2020) model learning effects that lead to long-term scarring after the pandemic is over as policy responses relating to debt forgiveness in the current pandemic can lead to lower leverage and consumption in the post-pandemic era. Through the learning channel in our model, there can also be “scarring” effects wherein agent’s consumption upon exit from a pandemic does not revert to the pre-pandemic levels due to the increase in updated probability of future pandemics. This, however, is not the focus of our paper.

Collin-Dufresne et al. (2016) show that learning can amplify the pricing of macroeconomic shocks when the representative agent has Epstein-Zin preferences and Bayesian updating. Our

⁸On a technical front, Hong et al. (2020b,a) consider aggregate transmission risk into SIR-style model, whereas our model of health risk arising from a pandemic is closer to the literature on rare disasters cited in the Introduction.

results on learning and the impact of parameter uncertainty on the value of a cure are related to the findings of both these papers; our model can generate both long-term scarring in consumption due to updated probability of pandemics and significant contraction of labor and consumption when parameter uncertainty is high, when the elasticity of intertemporal substitution (EIS) is low. Interestingly, expected time to deployment of a vaccine can be considered as a “macroeconomic shock” in our model that affects asset prices and depends crucially on parameter uncertainty in a manner that interacts with deep preference parameters.

3 Vaccine Progress Indicator and its Covariance with Stock Returns

As described in the introduction, the paper’s hypothesis is that the stock market may convey important information about the social value of resolving the pandemic. This section explains how we attempt to extract that information. There are two distinct steps. First, we construct a method for summarizing the state of vaccine research throughout 2020. Second, we estimate the stock market response to its changes.

3.1 Measuring Vaccine Progress

Readers are, by now, broadly familiar with the contours of the global effort to develop a vaccine for COVID-19. Through many of the excellent tracker apps, dashboards, and periodic survey articles we were all educated about the dozens of candidates under study, and their progress through pre-clinical work and clinical trials. On any given day, the state of the entire enterprise is a high-dimensional object consisting of multiple pieces of information about all of the projects. Our goal is to reduce that high dimensional object to a single number. Also, crucially, the number should have a tangible physical (or biological or economic) interpretation.

The single most salient aspect of vaccine development, the number that nearly all discussions boiled down to, was the anticipated time until widespread availability of a proven candidate. We therefore construct estimators of that quantity, using information that was available to observers at the time.

To do this, we introduce a stochastic model of candidate progress. 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). The model necessarily involves many param-

eters for which we have little hope of obtaining precise estimates. Details of our choices of all parameters are explained in Appendix B. We will validate our choices both by examining robustness to reasonable variation and by comparing them to other actual *ex ante* forecasts published during the sample period.⁹

Our stochastic model is a means to simulate, on each day t , the ultimate outcome of each of the candidates given their state of development as of that day. That simulated outcome, for candidate $n \in \{1, \dots, N\}$, is either reaching deployment by some date $T_n > t$ or not in which case $T_n = \infty$. Using the model, we can then run a large number, M , of joint simulations as of day t encompassing all of the candidates. In each simulation, m , the earliest time to widespread deployment is $T_m^* = \min_n \{T_n\}$. The cross-simulation average value of $T_m^* < \infty$ is the model's expectation T_t^D for the time to widespread deployment, conditional on at least one of the active candidates reaching deployment. Some fraction, μ , of simulations will result in all candidates failing and not reaching deployment. Our forecast is the full expectation, $\mathbb{E}_t T^*$ defined as $(1 - \mu)T_t^D + \mu T^{ND}$, where T^{ND} is an estimate of the time to deployment by a project other than those that are currently active (i.e., conditional on all active candidates not reaching deployment).¹⁰ In addition to the mean, the model also delivers the full distribution, and hence all quantiles, of T^* as of each date.

The model starts with a standard partition of the clinical trial sequences between pre-clinical, Phases I, II, III, application submission, and approval stages. A candidate in each stage either succeeds and moves to the next stage or fails.¹¹ We append a final stage: deployment, which an approved vaccine possibly still could not attain, e.g., due to technical infeasibility or emergence of serious safety concerns. Our assumption is that each stage is characterized by an unconditional probability of success, π_s , and an expected duration, τ_s .¹² We model each stage transition as a 2-state Markov chain with exponentially distributed times. Specifically, if we define two exponen-

⁹The appendix also presents evidence that our distributional assumptions are reasonably consistent with the (small) set of observed trial outcomes.

¹⁰The model does not attempt to forecast the entry of new projects.

¹¹This 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).

¹²Our success probabilities are taken from pharmaceutical research firm BioMedTracker and are based upon historical outcomes of infectious disease drug trials. Our duration estimates are based on projections from the pharmaceutical and financial press during 2020. See Appendix A for several examples.

tially distributed random variables t^u and t^d with intensities

$$\lambda_s^u = \frac{\pi_s}{\tau_s} \quad \text{and} \quad \lambda_s^d = \frac{1 - \pi_s}{\tau_s} \quad (1)$$

then it is straightforward to show that $t_s = \min\{t^u, t^d\}$ is exponentially distributed with intensity $1/\tau_s$ and that the probability of success, defined as $t^u < t^d$, is π_s .

Since our objective is to model the joint outcomes of all the vaccine projects, we need to specify the joint distribution of successes and failures. We do this by assuming the exponential times $\{t_n^u\}_{n=1}^N$ are generated by a Gaussian copula with correlation matrix \mathcal{R} , and likewise for the times $\{t_n^d\}$. The elements of the correlation matrix are set to positive values to capture the dependence between candidates. This positive dependence arises most obviously because all the vaccines are targeting the same pathogen, and will succeed or fail largely due to its inherent biological strengths and weaknesses. Beyond that, we also capture the dependence of candidates on one of a handful of strategies (or platforms) for stimulating immunological response. If, for example, an RNA-based platform proves to be safe and effective, then all candidates in this family would have a higher likelihood of success.¹³ Finally, some research teams (institutes or companies) have several candidate vaccines. Positive correlation between their outcomes may arise through reliance on common technological components, resources, or abilities.

As described so far, the “state” of a candidate is simply its current clinical-trial phase. This is not realistic in the sense that initiation of a new phase, as captured by the commencement of a new stage trial, is often known in advance. The trial start date (as reported to the U.S. government) may not actually be the date of the arrival of news about the candidate. Likewise, within a stage (as a trial is progressing), the “state” is not static. Information about the trial (preliminary results), or more complete information about earlier trials, may be published or released to the press or leaked. Trial schedule information (delays or acceleration) may be announced. Regulatory actions by non-U.S. authorities may also convey information. For all of these reasons, we modify our framework by adjusting the probability of each candidate’s current-stage success on the date of arrival of news specific to it.¹⁴ Because positive news is more likely to be revealed than negative

¹³In October, two candidate vaccines had their trials paused due to adverse reactions: both were based on adenovirus platforms.

¹⁴Our source of news is FactSet StreetAccount. We classify vaccine related stories into seven positive types and six negative types. The types and probability adjustments are given in the appendix.

news, we also deterministically depreciate each candidate’s success probability in the absence of news. We will verify below that our conclusions about the stock market response to vaccine progress are not driven by our assumptions regarding the arrival of interim trial news.¹⁵

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 these countries.¹⁶ This assumption is consistent with the progress of these candidates receiving minimal coverage in the U.S. financial press. We will also verify that including them in our index does not change our primary results.

It is worth acknowledging that, in focusing on the scientific advancement of the individual candidates, our measure does not currently attempt to capture general news about the vaccine development environment and policy. For example, 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. The Fall of 2020 saw open debate about the standards that would be applied for regulatory approval, the outcome of which could have affected forecasts as well. While we could alter our index based on some subjective 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, and Figure 2 shows the number of active vaccine projects. 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 declin-

¹⁵Technically, altering the marginal success probabilities within a trial induces a non-exponential unconditional marginal distribution of trial duration. We retain the exponential assumption of the Gaussian copula for tractability. Our results are robust to using constant probabilities.

¹⁶We retain candidates coded as multi-country projects including Russia or China.

ing, 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.

3.1.1 Validation

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

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 each bin. For each survey we estimate the median response, assuming a uniform distribution of responses within the bin containing the median.¹⁷ 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 Judgment 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 4 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 reason is the specificity of the particular survey question, which specifies an exact quantity of the vaccine being distributed in

¹⁷While it is tempting to equate the surveys’ distribution of forecasts with a forecasted distribution, these are conceptually distinct objects that need not coincide. In addition, in each survey, the farthest out forecast bin is unbounded, meaning that “never” (or “more than 3 years from now”) is a possible response. So, for both reasons, it is problematic to compute a weighted average forecast across the response bins. The modal response bin is also not a good summary statistic for the same reason, and also because it depends on the bin widths.

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.

3.2 Stock Market Response

Figure 3 plots vaccine progress (inverted) along with the market portfolio's year-to-date performance. In principle, assessing the stock market response to changes in vaccine progress should be straightforward. However, the circumstance of 2020 complicate the task. In a nutshell, there was a lot else going on. The amount of information for markets to digest was enormous and multifaceted. Even the information flow just about coronavirus research *other than vaccine trials* was voluminous. Thus, how to control for non-vaccine related news becomes an important consideration.

Our approach is to run daily regressions of stock market returns on vaccine progress and exclude days that contained large stock market moves that have been reliably judged to have been due to other sources of news. Specifically, we employ the classification of Baker et al. (2020b) for causes of 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, research on vaccines falls under their "other" category, whereas news about the pandemic itself was usually categorized as "macroeconomic". We view market returns on such days as very unlikely to have been driven by vaccine news if none of the three analysts assigns more than 25% weight to the other 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,¹⁸ and on the prior knowledge that positive vaccine progress cannot be negative news. We include dummies for all of the non-vaccine large-news days. There are 28 such days, 17 of which were in March.

Our approach is imperfect. We have no other controls outside these large move days when there were certainly other factors influencing markets. Including dummy variables effectively reduces our sample size. However, at a minimum we are limiting the ability of our estimation to misattribute the largest market moves to vaccine progress.

¹⁸As of the time of this draft, Baker et al. (2020b)'s website had classified days through June. We append September 3 and September 23 as two dates with negative jumps but arguably were driven by non-vaccine progress related news.

Table 2 shows the resulting regression estimates of market impact. These regression specifications include changes in the vaccine progress indicator in a five day window around each day, t , on which stock returns are measured. Including changes on days other than the event day- t guards against our imperfect attribution of the date of news arrival. *A priori* we suspect it is more likely that, if anything, markets have information before it is reflected in our index, meaning the relevant reaction would correspond to the $t + 1$ or $t + 2$ coefficients. On the other hand, 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 $t - 1$ or $t - 2$ coefficients. The specifications also include two lags of the dependent variable to control for short-term liquidity effects. Specifically, the regression is

$$R_{m,t}^e = \alpha + \sum_{h=-2}^2 \beta_h \Delta \text{VPI}_{t+h} + \gamma_1 R_{m,t-1}^e + \gamma_2 R_{m,t-2}^e + \sum_{j=1}^{28} \delta_j \mathbb{1}_{\text{jump } j} + \epsilon_t \quad (2)$$

where ΔVPI_t is the change in vaccine progress indicator, and $\mathbb{1}_{\text{jump } j}$ is a dummy equal to one on the j th jump date from Baker et al. (2020b). The dependent variable is the return to the value-weighted CRSP index from January 1 through October 31, 2020.

The first column of the table shows results using our baseline vaccine progress indicator. The coefficient pattern shows the largest negative responses on the $t - 1$ and $t + 2$ index changes. Focusing on the cumulative impact, the sum of the β s is statistically significant at the 1% level. The precisely estimated point estimate implies a stock market increase of 8.6 percent on a decrease in expected time to deployment of one year. This number seems plausible: subsequent to our sample, on November 9th the U.S. stock market opened almost 4% higher in response to positive news from Pfizer on Phase III trial results. This would imply more sensitivity than the OLS estimate if, as seems likely, the news revised estimates of time to deployment by less than six months.¹⁹

Returning to Table 2, the second and third columns implement the methodology of Kogan et al. (2017) (hereafter KPSS). Those authors use an empirical Bayes procedure to estimate the market value of patents using the stock returns to the patenting firm in an event window sur-

¹⁹While it is not the focus of the paper, it is also interesting to ask about the total contribution of vaccine progress to the stock market performance during the sample period, and to the post-March rally in particular. From March 23 to October 30 our forecast dropped by 2.5 years, of which 0.6 years was expected. The OLS point estimate then implies that vaccine news in total would have induced a 16.3% positive return. The return on the S&P 500 during this period was 47.7%. Hence, vaccine progress could have been responsible for 34% (16.3/47.7-1) of the rally.

rounding patent publication date. As in our case, economic logic rules out a negative response: vaccine progress cannot be unfavorable news just as the value of a patent must be positive. KPSS employ a truncated normal prior distribution for the unobserved true response. Conditional on knowing the return standard deviation, the posterior mean estimate of the response coefficient is then also distributed as a truncated normal. The estimation methodology generalizes naturally to a multivariate regression setting²⁰ (O’hagan (1973)). The table reports posterior mean and standard deviations for the individual response coefficients and for their sum.²¹ The methodology is sensitive to the specification of the prior variance of the coefficient distribution. Both column 2 and column 3 assume that the pre-truncated normal distribution for β_t has standard deviation equal to 1, which, after truncation, implies that 84% of the distribution mass is below 1.0. We regard this as a conservative (or skeptical) choice.²² Results in the second column use the same (independent) prior for all the response coefficients. The third column uses a smaller prior mean for the lead and lag coefficients.²³ Both priors produce posterior means for the sum of the five response coefficients that are lower than the OLS estimate: -6.4% in column 2 and -4.1% in column 3. Note that the estimation is sharp in both cases in the sense that each posterior mean is several standard deviations from zero. The calibrations in the next section will adopt the range of these conservative estimates.

To examine the robustness of the response estimates to the assumptions built into the vaccine progress index, we repeat the OLS specification estimation with five variants. These results are shown in Table 3. The first column repeats the original specification from the prior table. The next two columns vary the assumptions about the effect of news to phase success probabilities. (Column 2 includes no news adjustments. Column 3 applies the news adjustments to only the current trial phase, as opposed to all future phases, and increases the $\Delta\pi$ from news on positive data releases, positive enrollment and dose starts to 15%, 5% and 5%, respectively) The fourth column increases the base copula correlation from 0.2 to 0.4. The fifth column lowers the assumed prob-

²⁰We follow KPSS in assuming a zero mean under the prior for the pre-truncated normal distribution, assuming returns are normally distributed, and in using the regression residual to estimate the return standard deviation. Note that the estimation still includes dummy variable for market jump days making the normality assumption plausible.

²¹Moments of the truncated multivariate normal posterior are computed using the algorithm of Kan and Robotti (2017) using software provided by Raymond Kan. <http://www-2.rotman.utoronto.ca/~kan/research.htm>.

²²Note that making the prior more diffuse does not, in this case, correspond to making it less informative: the prior mean increases with the standard deviation.

²³Specifically, the assumption is that pre-truncated standard deviations are 0.7 for the first lead and lag and 0.5 for the second lead and lag.

ability of successful deployment following approval. Finally the sixth column includes vaccine candidates whose research program is based in Russia or China. In all of these cases the sum of the response coefficients is highly statistically significant and the point estimates are in the same range as those in Table 2.²⁴

3.3 Industry Responses

As a validity check for our primary 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, 2020²⁵ and a national emergency declared on March 13, 2020.²⁶ 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, 2020²⁷ and significant advances in vaccine progress, helping us pin down industry covariances with COVID-19 itself, separate from covariances with monetary policy responses and vaccine progress.

We then estimate industry sensitivity to vaccine progress over the non-overlapping sample from March 23, 2020 to October 31, 2020. Specifically, we re-estimate (2) sector-by-sector,

$$R_{i,t}^e = \alpha + \sum_{h=-2}^2 \beta_{h,i} \Delta \text{VPI}_{t+h} + \gamma_{1,i} R_{i,t-1}^e + \gamma_{2,i} R_{i,t-2}^e + \sum_{j=1}^{28} \delta_{j,i} \mathbb{1}_{\text{jump } j} + \epsilon_{i,t} \quad (3)$$

where $R_{i,t}^e$ denotes value-weighted excess returns on the 49 Fama-French industry portfolios.

Figure 4 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

²⁴We further explore adjusting each candidate’s state-level duration, in addition to probability of success, using relevant news. Our findings are robust to these additional specifications. Results are available upon request.

²⁵<https://www.hhs.gov/about/news/2020/01/31/secretary-azar-declares-public-health-emergency-us-2019-novel-coronavirus.html>

²⁶<https://www.whitehouse.gov/presidential-actions/proclamation-declaring-national-emergency-concerning-novel-coronavirus-disease-covid-19-outbreak/>

²⁷<https://www.federalreserve.gov/monetarypolicy/pmccf.htm>

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. Hence, the results here make it unlikely that our primary findings on the market’s sensitivity are due to omitted variables.

4 Regime-Switching Model of Pandemics

In this section, we introduce a regime-switching model of pandemics in order to derive the value of a cure in terms of the economy’s primitive objects, such as the ratio of labor supply or marginal propensity to consume between the pandemic and the non-pandemic states. In order to connect the theory to our empirical exercise, we need a model with four attributes: a description of pandemics; a well-defined notion of the value of ending a pandemic; a depiction of progress towards that objective; and a stock market that is sensitive to that progress.

4.1 S-State Model

We consider the state of the economy to be either in “non-pandemic” regime or in “pandemic” regime.²⁸ Within the pandemic regime, there can be several sub-states that correspond in our context to different stages in the development of vaccines. We denote the state as $s \in \{0, 1, \dots, S - 1, S\}$, where for ease of notation both state 0 and state S are the same non-pandemic states, and the others are pandemic states. We assume that the economy switches between these states based on a Markov-switching or transition matrix. The transition probabilities are given as follows where η , the probability of switching from the non-pandemic regime to the pandemic regime, and λ_d and λ_u , the respective probabilities in a pandemic state to move “down” or “up” to the adjacent states,

²⁸In the Appendix, we work out in detail the solution to the 2-state regime-switching model in which the pandemic regime consists of just one state. Besides illustrating the detailed solution to the model (Hamilton-Jacobi-Bellman (HJB) equations, labor and consumption choices, and system to determine the value function), it also serves as the benchmark case for developing the model further with parameter uncertainty.

are assumed to be exogenous:

$$P(s_{t+dt}=1|s_t=0 \text{ or } S) = \eta dt \quad (4)$$

$$P(s_{t+dt}=s_t|s_t=0 \text{ or } S) = 1 - \eta dt \quad (5)$$

$$P(s_{t+dt}=s-1|s_t=s \in [1, S-1]) = \lambda_d(s) dt \quad (6)$$

$$P(s_{t+dt}=s+1|s_t=s \in [1, S-1]) = \lambda_u(s) dt \quad (7)$$

$$P(s_{t+dt}=s_t|s_t=s \in [1, S-1]) = 1 - \lambda_d(s) dt - \lambda_u(s) dt \quad (8)$$

4.1.1 Agents, Labor and Capital Stock

We assume the economy has a unit mass of identical agents, each with the following production function that in the non-pandemic state depends on the labor input l and generates a stochastic output q gross of consumption as:

$$l^\alpha q \mu dt + \sigma l^{\alpha/2} q dB_t \quad (9)$$

where $\{B_t, t > 0\}$ is a Brownian Motion process. We can view q as capital stock – physical and human – or wealth of the agent that is readily convertible into consumption (a form of “buffer stock” therefore), and $\alpha \in (0, 1)$ is the elasticity of instantaneous expected output with respect to labor. The instantaneous expected return is $l^\alpha \mu dt$ and the instantaneous variance in output is $l^\alpha \sigma^2 dt$. We will assume $l \in [0, \bar{l}]$, where \bar{l} is an upper bound representing technological or human constraints on investment into the capital stock.

The production function of agent in the non-pandemic state, net of consumption flow, is therefore:

$$dq = l^\alpha q \mu dt - C dt + \sigma l^{\alpha/2} q dB_t \quad (10)$$

where C is the endogenous consumption rate.

In the pandemic regime, the production function has all of the above features but it becomes exposed to the risk of an economic loss when hit by a “health” disaster:

$$dq = l^\alpha q \mu dt - C dt + \sigma l^{\alpha/2} q dB_t - [l\varepsilon + k + KL] q dP(t). \quad (11)$$

Then let

$$\chi(l, L) \equiv [l\varepsilon + k + KL], \quad (12)$$

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 to the pandemic via aggregate labor. $P(t)$ is a Poisson process with a parameter ζ and a known jump amplitude $\Delta \in (0, 1)$. When the Poisson process is triggered, a part of the agent's capital stock is destroyed and falls to $q(1 - \chi\Delta)$, e.g., due to health-induced disruptions to work, the need to work from home with attendant productivity impact and loss of human capital, filing for bankruptcy with deadweight loss in asset value, and firing of labor with difficulty to re-match at a future date, etc. We will assume parametric restrictions on ε , k and K to be small enough that $(1 - \chi\Delta) \in (0, 1)$. The specification allows for both the labor supply choice (l) and the pre-existing conditions of the household unrelated to labor supply (k) to potentially amplify the health shocks. In addition, aggregate labor supply (L) can also amplify exposure to the pandemic as a form of externality. The agent takes the aggregate supply of labor L as given in her optimization problem.

4.1.2 Agent's Preferences

We assume that 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

$$\mathbb{J}_t = \mathbb{E}_t \left[\int_t^\infty f(C_{t'}, \mathbb{J}_{t'}) dt' \right] \quad (13)$$

and aggregator

$$f(C, \mathbb{J}) = \frac{\rho}{1 - \psi^{-1}} \left[\frac{C^{1-\psi^{-1}} - [(1-\gamma)\mathbb{J}]^{\frac{1}{\theta}}}{[(1-\gamma)\mathbb{J}]^{\frac{1}{\theta}-1}} \right] \quad (14)$$

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

$$\theta^{-1} \equiv \frac{1 - \psi^{-1}}{1 - \gamma} \quad (15)$$

We also assume that the state of the economy s is known to each agent and so are the transition probabilities. Later on, we will consider parameter uncertainty for a two-state version of the model.

4.1.3 Agent's Optimization Problem and Equilibrium

The representative agent's problem is to choose in each state s optimal consumption $C(s, L^*(s))$ and labor $l(s, L^*(s))$ that maximizes the objective function $\mathbb{J}(s)$; in particular, the agent must 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:

$$dq(s) = [l(s, L^*(s))]^\alpha q \mu dt - C(s, L^*(s)) dt + \sigma [l(s, L^*(s))]^{\alpha/2} q dB - \chi(l(s, L^*(s)), L^*(s)) q dP(t) \quad (16)$$

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

4.1.4 Solution

We show in the Appendix how to set up the HJB equation for the optimization problem of the representative agent to determine the optimal consumption $C(s, L^*(s))$ and labor $l(s, L^*(s))$, making the natural conjecture that the value function is given by

$$\mathbb{J}(s) \equiv \frac{H(s)q^{1-\gamma}}{1-\gamma} \quad (17)$$

where $H(s)$ are constants (depending on deeper parameters of the model) to be determined. Given the transitions across states, $H(s)$ are jointly determined as explained below; however, given the structure of the problem, the equilibrium labor choices are more simply derived.

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

$$L(0) = L(S) = \bar{\ell} \quad (18)$$

Equilibrium labor in pandemic states $L^*(s) \forall s \in \{1, \dots, S-1\}$ solves²⁹

$$\chi(L(s), L(s)) = k + (\varepsilon + K)L(s) = \frac{1}{\Delta} \left[1 - (L(s))^{\frac{1-\alpha}{\gamma}} \nu \right] \quad (20)$$

where

$$\nu \equiv \left[\frac{\alpha \left(\mu - \frac{1}{2} \gamma \sigma^2 \right)}{\zeta \varepsilon \Delta} \right]^{-\frac{1}{\gamma}}. \quad (21)$$

Note: All proofs appear in the appendix.

In the non-pandemic state, the agent faces no cost to supplying labor to augment the capital stock 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 stock that arises when the pandemic hits. A key property of the model is that as the exposure to the pandemic is a function of the labor supply, the agent contracts labor in general relative to the non-pandemic state. We will assume parameter restrictions are such that this is in fact the case.

Next, the constants $H(s)$ that pins down the agent's equilibrium objective function in state s are solved jointly as follows:

Proposition 2. Denote

$$g(x, y) \equiv \frac{(1-\gamma)\rho}{(1-\psi^{-1})} - x^\alpha (1-\gamma) \left(\mu - \frac{1}{2} \gamma \sigma^2 \right) - y \left([1 - \chi(x, x)\Delta]^{1-\gamma} - 1 \right) \quad (22)$$

²⁹It can be shown that given $\alpha \in (0, 1)$, the second order condition for a maximum is satisfied whenever

$$\mu - \frac{1}{2} \gamma \sigma^2 > 0 \quad (19)$$

which also implies $\nu > 0$.

Then $\{H\}$'s are given by the system of S recursive equations $E(0), \dots, E(S-1)$ as follows:

$$E(0) : g_0 \equiv g(\bar{\ell}, 0) = \frac{(1-\gamma)}{(\psi-1)} \rho^\psi (H(0))^{-\psi\theta-1} + \eta \left[\frac{H(1)}{H(0)} - 1 \right] \quad (23)$$

$$E(s) : g_1 \equiv g(L(s), \zeta) = \frac{(1-\gamma)}{(\psi-1)} \rho^\psi (H(s))^{-\psi\theta-1} + \lambda_d(s) \left[\frac{H(s-1)}{H(s)} - 1 \right] + \lambda_u(s) \left[\frac{H(s+1)}{H(s)} - 1 \right] \quad (24)$$

for $s \in \{1, \dots, S-1\}$, and $H(S) = H(0)$.

This captures another key property of the model. The solution to determining the agent's objective functions depends crucially on the relative values of g_0 and g_1 , which serve as an important statistic for the extent of labor contraction in the economy in pandemic states relative to the non-pandemic state. The lower is g_1 relative to g_0 , the lower is the objective function in pandemic states relative to the non-pandemic state, and in turn, the greater is the value that the agent attaches to finding a cure that can effect a switch out of the pandemic.

Corollary 1. *Pandemic regime parameters $(\zeta, \varepsilon, \kappa, K, \Delta)$ affect the agent's objective function, i.e., $H(s)$, only through the value of g_1 .*

This observation will play a crucial role in using our empirical work to infer the value of a cure. Next,

Proposition 3. *Equilibrium consumption in state s can be determined based on $H(s)$ as*

$$C(s) = \frac{(H(s))^{-\psi\theta-1} q}{\rho^{-\psi}} \quad (25)$$

Marginal propensity to consume, $c(s)$, which is defined as dC/dq , depends on the state (s) and the elasticity of intertemporal substitution (ψ). Figure 5 illustrates for a 10-state regime-switching model that $c(s)$ in the pandemic states, $s \in \{1, \dots, 9\}$, is below (above) that in the non-pandemic state, $s = 0$ or $s = 10$, when ψ is below (above) 1.

Combining Propositions 2 and 3, the equilibrium can also be characterized in terms of labor and consumption outcomes in different states, s , the solution to the following system of recursive equations,

Corollary 2. *Following Proposition 3, we can write the system of S recursive equations $\hat{E}(0), \dots, \hat{E}(S-1)$ that characterize the $\{c\}$'s as:*

$$\hat{E}(0) : g_0 \equiv g(\bar{\ell}, 0) = \frac{(1-\gamma)}{(\psi-1)}c(0) + \eta \left[\left(\frac{c(1)}{c(0)} \right)^{-\theta\psi^{-1}} - 1 \right] \quad (26)$$

$$\hat{E}(s) : g_1 \equiv g(L(s), \zeta) = \frac{(1-\gamma)}{(\psi-1)}c(s) + \lambda_d(s) \left[\left(\frac{c(s-1)}{c(s)} \right)^{-\theta\psi^{-1}} - 1 \right] + \lambda_u(s) \left[\left(\frac{c(s+1)}{c(s)} \right)^{-\theta\psi^{-1}} - 1 \right] \quad (27)$$

for $s \in \{1, \dots, S-1\}$, and $c(S) = c(0)$.

4.1.5 Externality and the Central Planner

Before proceeding to the value of a vaccine in this setup, it is worth noting that in our model there is an externality that the impact of labor on the effect of the pandemic via KL term (where L is the aggregate labor) is not internalized by each agent. The planner would factor this in the socially efficient choice of labor for each agent. This is tantamount to replacing ε by $(\varepsilon + K)$ in ν above to obtain ν^{CP} :

$$\nu^{CP} \equiv \left[\frac{\alpha \left(\mu - \frac{1}{2} \gamma \sigma^2 \right)}{\zeta (\varepsilon + K) \Delta} \right]^{-\frac{1}{\gamma}} \quad (28)$$

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

$$\chi(L(s), L(s)) = k + (\varepsilon + K)L(s) = \frac{1}{\Delta} \left[1 - (L(s))^{\frac{1-\alpha}{\gamma}} \nu^{CP} \right] \quad (29)$$

It is then straightforward to show that $\nu^{CP} > \nu$ for $K > 0$ and $\gamma > 0$, and $L^{CP}(s) < L(s)$, i.e., the socially efficient choice of labor in pandemic states is smaller than the privately optimal one. We will see later that this insight will help understand the wedge between the values attached to a cure by the planner and the private agents.

4.1.6 Value of a Cure

We have now all the ingredients to determine the value of a cure. We define it as 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):

$$V(s) = 1 - \left(\frac{H(s)}{H(0)} \right)^{\frac{1}{1-\gamma}} \quad (30)$$

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-pandemic state.³⁰

Using the optimal consumption characterized above, we obtain that

Proposition 4. *The value of a cure in the pandemic state s is determined by the ratio of marginal propensity to consume ($c \equiv dC/dq$) in the pandemic state s relative to that in the non-pandemic state, adjusted by the agent's elasticity of intertemporal substitution (EIS):*

$$V(s) = 1 - \left(\frac{c(s)}{c(0)} \right)^{-\frac{1}{\psi-1}} = 1 - \left(\frac{C(s)}{C(0)} \right)^{-\frac{1}{\psi-1}} \quad (31)$$

Note from Proposition 3 that the ratio of marginal propensities to consume for a given q is simply the ratio of consumptions. Furthermore, when EIS (ψ) is less than one, the value of a cure is higher the greater is the contraction of consumption in the pandemic states relative to the non-pandemic state. Indeed, it can be shown that consumption is lower in the pandemic states relative to the non-pandemic state (for a given q) if and only if $\psi < 1$.

Our next step is to derive this value. To this end, we will derive asset prices in the framework above to show that the value of claim to the economy's output ("stock market") changes in relation to the expected time to switching out of a pandemic ("expected time to deployment of a vaccine") – which we empirically estimated – is crucially determined by the contraction of labor in pandemic state relative to the non-pandemic state, as described by $\frac{g_1}{g_0}$. This helps us estimate the value of a cure under assumptions of standard preference parameters.

4.2 Connection to the Data

We introduced the model in order to first define and study determinants of welfare and the value of a vaccine. The second reason to introduce the model is to allow us to bring financial claims

³⁰We acknowledge that this is essentially a comparative static exercise and the economy does not possess the technology to actually effect this transition. We discuss in the Conclusion ways to enrich our model to introduce the vaccine technology into the model as an important topic for future research, but one that is beyond the immediate scope of this paper.

into the picture, and, in particular, to examine the model’s counterpart to the sensitivity that we estimated in Section 2.2.

As is standard in the asset pricing literature, we begin by interpreting “the market portfolio” within the model as a claim to the economy’s output.³¹ Output is the net new resources per unit time, which is implicitly defined by two endogenous quantities: the change in the cumulative physical capital plus consumption, or $dq + Cdt$. (Note that the value of a claim to this flow is not equal to the value of a claim to the capital stock, which is q .) 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. From this, we derive the value of the claim and some key properties in the following proposition.

Proposition 5. *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. The system depends on the pandemic parameters through only two quantities, which may be taken to be the risk-neutral expected growth of output and g_1 , defined in the preceding section.*

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 .

Next, define T^* as the time at which the state S is attained and the pandemic is terminated. Assuming the progression and regression intensities λ_u and λ_d are constant, it is straightforward to show that its time t expectation, $\mathbb{E}_t[T^*]$ is again given by a linear system, which we omit for brevity. Moreover, for large S , the difference

$$\Delta \mathbb{E}[T^*] = \mathbb{E}[T^* | s + 1] - \mathbb{E}[T^* | s] \sim \frac{1}{\lambda_u} \quad (32)$$

is effectively constant as well. (The expression is exact for $s = 1$.)

Combining the above two results, we can readily define the model’s analogue of the sensitivity

³¹We defer for now explicitly modeling the dividend share of output or incorporating leverage in the equity claim.

that we empirically estimated as

$$\frac{\Delta \log P}{\Delta \mathbb{E}[T^*]} \quad (33)$$

For our purposes the crucial property of this quantity is that it allows us to approximately pin down the pandemic parameters that determine the value of a cure. Specifically, it depends importantly on g_1 , modulated by the other pandemic parameters. This is illustrated in Figure 6, which plots the sensitivity for a wide range of model solutions. The horizontal axis is g_1 and each model version corresponds to a single point. Here we allow all of the pandemic parameters to vary, as well as the intensities η and λ_u . From the proposition above, g_1 alone does not suffice to determine the pricing function p . However, the second variable that the proposition identifies as mattering – the risk neutral expected growth of output – is codetermined in equilibrium with g_1 , and as a practical matter, its residual variation is small. That is, given g_1 , the expected growth rate varies only marginally with the remaining parameters.

Hence, although the identification is not exact, we can infer from the figure that our empirical estimates in the range of 5.0 are consistent with a value of g_1 in range of approximately -0.39 to -0.37, given the non-pandemic parameters used to compute the model solutions shown.

4.3 Calibrating the Value of a Cure

In this section and in Section 5 we present comparative static results exploring the determinants of the value of a cure, V as defined in Section 4.1. In doing this, unless otherwise stated, we will fix the non-pandemic parameters to be the values shown in Table 4. The preference parameters are broadly consistent with the asset pricing literature under stochastic differential utility, although we use a relatively low level of risk aversion because higher values of γ can lead to violations of regularity conditions. The growth rate and volatility parameters are chosen as a compromise between two interpretations of the model. On the one hand, we are viewing the output process as representing national income (or GDP), which would suggest smaller mean and volatility. On the other hand, our asset pricing exercise views the same process as depicting dividends, which would suggest higher values for both.³² In addition, the solutions in this section will set the

³²An additional consideration is that a relatively high growth rate is needed to obtain a solution when varying the elasticity of intertemporal substitution, which we will do in Section 4.

number of states to be $S = 10$, which is arbitrary but without loss of generality. Our results are not too sensitive to the specific choice of the number of states as it is to the other pandemic parameters. We also set the intensity of regress to be $\lambda_d = 0$, which limits vaccine related volatility. This choice accords with recent experience: the research setbacks through the Fall of 2020 were few and had little impact on our measure of progress.

Figure 7 plots the value of the cure as a function of the remaining timing parameters, η and λ_u (hereafter we will denote $\lambda_u / (S + 1)$ as λ without a subscript). The left panel plots V against $1/\lambda$ the expected duration of the pandemic, while the right panel uses the pandemic frequency η on the horizontal axis. (The left panel sets $\eta = 0.02$ and the right panel sets $\lambda = 0.5$. Both panels take the current state as $s = 1$.) From the left plot, agents in the economy would be willing to give up five percent of their wealth for an immediate transition to state 0 even when the pandemic is only expected to last one year. This value rises to approximately 15% when the expected duration is 4 years. The right panel shows that the value of a cure is actually lower when pandemics are more frequent. Recall that a “cure” here only applies to the current pandemic. A one-time cure is less valuable when a new one will be needed sooner.

Given the parameters used in the calibration, the (endogenous) expected decline in wealth due to pandemic shocks is approximately 5% per year. Our estimation of the value of a cure is quite close to this expected loss, which is intuitively sensible. Also, while the two quantities are conceptually distinct, the value we are computing here is similar (on a per year basis) to the stock market valuation of a year of pandemic experience as estimated in Section 3.

Table 5 shows the effect on V of the labor market externality, for a range of λ and η . Here the right panel shows the benchmark case while the left panel shows what happens when the labor market response to the pandemic is determined by a welfare maximizing central planner. The result shows a small but not insignificant increase in the value of the vaccine in the presence of the externality. In effect, the extra degree of lock-down that the planner would impose and the vaccine are substitutes as countermeasures. Note that our baseline parameters in Table 5 imply a labor contraction of approximately 25% in the pandemic regimes relative to the non-pandemic

regime, in line with observed empirical data.^{33,34} Before proceeding, 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 vaccine more than the representative agent.

5 Learning and Uncertainty

We have used the S -state version of our model to study the reaction of markets to vaccine news within a pandemic. Relating its predictions to the empirical evidence in Section 3 has provided evidence on plausible parameters affecting the value of a vaccine. Now we return to the two-state version of our model in order to examine the role of vaccine news from a different angle. Specifically, we are interested in the accumulation of information over longer horizons about the frequency and duration of pandemics. We study the effect upon the value of a vaccine of uncertainty about these quantities and of differing attitudes towards uncertainty.

5.1 Information Structure

Recall that in the two-state model η is the intensity of switching from state 0 (“off”) to state 1 (“on”) and λ is the intensity of switching from 1 to 0. In this section, we assume that agents have imperfect information about these intensities.

Let us stipulate that at time zero the agent has beliefs about the two parameters that are described by gamma distributions, which are independent of each other. Each gamma distribution has a pair of non-negative hyperparameters, a^η, b^η and a^λ, b^λ , that are related to the first and sec-

³³Muellbauer (2020) models a larger drop in consumption than income during the pandemic with a credit-augmented consumption function. Using customized survey data, Coibion et al. (2020a,b) find the pandemic led to a 20 million decline in the number of employed workers by the first week of April 2020, and attributed 60 percent of the decline in the employment-to-population ratio by May 2020 to lockdowns. Dingel and Neiman (2020), Mongey et al. (2020) and Beland et al. (2020) classify occupations by their work from home feasibility, documenting more adverse labor market outcomes for occupations with high proximity among coworkers. For those looking for employment, Forsythe et al. (2020) find job vacancies had fallen 40% by April 2020 compared to pre-COVID-19 levels, with the largest declines in leisure, hospitality and non-essential retail. Consequently, Bernstein et al. (2020) find a flight-to-safety effect, with job seekers shifting searches from early-stage ventures to larger firms, while also considering lower salaries, and alternative roles and locations.

³⁴Baker et al. (2020a) deploy transaction-level data to study consumption responses to the pandemic, 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. (2020a) 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.

ond moments via

$$\mathbb{E}[\eta] = \frac{a^\eta}{b^\eta}, \quad \text{Std}[\eta] = \frac{\sqrt{a^\eta}}{b^\eta}, \quad (34)$$

and likewise for λ .

By Bayes' rule, under this specification, as the agent observes the switches from one regime to the next, her beliefs remain in the gamma class with the hyperparameters updating as follows

$$\begin{aligned} a_t^\eta &= a_0^\eta + N_t^\eta \\ b_t^\eta &= b_0^\eta + t^\eta \end{aligned}$$

where t^η represents the cumulative time spent in state 0 and N_t^η represents the total number of observed switches from 0 to 1. Analogous expressions apply for λ . Thus, during the "off" regime, the only information that arrives (on a given day, say) is whether or not we have switched to "on" on that day. If that has occurred, the counter N^η increments by one and the clock t^η turns off (and t^λ turns on). In this version of the model, that is the entirety of the information revelation. In contrast to the previous section, no good or bad news arrives about progress during a regime. Although this setting lessens the model's ability to speak to high-frequency dynamics, it allows us to study the role of uncertainty in the economy's longer term evolution.

Under the above information structure, the economy is characterized by a six-dimensional state vector consisting of the stock of wealth, q , $a^\eta, b^\eta, a^\lambda, b^\lambda$ and the regime indicator S . However this six-dimensional space can actually be reduced to three.

Since the switches between states alternate, let us define an integer index M_t to be the total number of switches $N_t^\eta + N_t^\lambda$ and then (assuming we are in state 0 at time 0) $N_t^\eta = M_t/2$ when M is even, and $N_t^\lambda = (M_t + 1)/2$ when M is odd. Knowing M (along with the priors a_0^η and a_0^λ) is equivalent to knowing a_t^η and a_t^λ . Given these values, specifying the current estimates

$$\hat{\eta}_t \equiv \mathbb{E}_t[\eta] \quad \text{and} \quad \hat{\lambda}_t \equiv \mathbb{E}_t[\lambda] \quad (35)$$

is equivalent to specifying the remaining hyperparameters b_t^η and b_t^λ . Thus, solutions to the model can be described as a sequence of functions $H_M(\hat{\eta}, \hat{\lambda})$ for the agent's value function at step M .

Compared to the full-information model in Section 4, within each regime the only new changes to the state come through variation in the estimates $\hat{\eta}_t$ and $\hat{\lambda}_t$ which change deterministically with the respective clocks t^η and t^λ . Holding M fixed, the dynamics of $\hat{\eta}_t$ are given by

$$d\hat{\eta}_t = d\frac{a_t^\eta}{b_t^\eta} = a_t^\eta d\frac{1}{b_t^\eta} \quad (36)$$

$$= -\frac{a_t^\eta}{(b_t^\eta)^2} dt \quad (37)$$

$$= -\frac{(\hat{\eta}_t)^2}{a_t^\eta} dt. \quad (38)$$

Under partial information, we proceed as in Section 4 to write-out the HJB equation with the state variables following the dynamics determined by the representative agent's information set. As before, we can conjecture a form of the value function

$$\mathbb{J} = \frac{q^{1-\gamma}}{1-\gamma} H(\hat{\eta}, \hat{\lambda}, M; C, \ell). \quad (39)$$

And, as before the first order condition for consumption yields $C = q (\rho^\psi) H_1^e$ (where e_1 is defined in Section 4.1). This follows because consumption does not enter into any of the new terms involving the information variables. Also fortunately, none of the information variables appears in terms affected by labor supply, ℓ , and the function H drops out of the first-order condition for ℓ . (Intuitively, nothing about the likelihood of changing regimes affects the optimal choice of labor within a regime.) This means that the solutions for ℓ^* can be computed independent of the rest of the system.

Using these the results, the HJB system can be written as the infinite-dimensional linked PDEs:

$$g_0 = \rho^\psi \left(\frac{\theta}{\psi}\right) H_M^{-\psi/\theta} + \hat{\eta} \left(\frac{H_{M+1}}{H_M} - 1\right) - \frac{(\hat{\eta})^2}{a^\eta H_M} \frac{\partial H_M}{\partial \hat{\eta}} \quad (40)$$

$$g_1 = \rho^\psi \left(\frac{\theta}{\psi}\right) H_{M+1}^{-\psi/\theta} + \hat{\lambda} \left(\frac{H_{M+2}}{H_{M+1}} - 1\right) - \frac{(\hat{\lambda})^2}{a^\lambda H_{M+1}} \frac{\partial H_{M+1}}{\partial \hat{\lambda}} \quad (41)$$

where M runs over the even integers.³⁵

For large M , the estimation errors for both η and λ , expressed as a fraction of the posterior

³⁵The constants g_0 and g_1 are as defined in Section 4. See the internet appendix for a derivation of (40)-(41).

estimates, go to zero:

$$\frac{\text{Std}[\eta]}{\mathbb{E}[\eta]} = \frac{1}{\sqrt{a^\eta}} = \frac{1}{\sqrt{a_0^\eta + M_t}}. \quad (42)$$

Hence the system always converges to the full-information solution. This provides a boundary condition, which, together with the single-regime solutions on the edges of the $(\hat{\eta}, \hat{\lambda})$ plane, enables computation of all the individual H functions.³⁶ It can be shown that, as in the full-information case, a necessary and sufficient condition for existence of a solution is $g_0 > g_1$.

As in the previous section, once the value function is obtained, we can characterize the certainty equivalent value of a vaccine that produces an immediate transition from the pandemic state to the non-pandemic state. The next section performs this calculation and analyzes the drivers of variation in that value.

5.2 Results

Table 6 shows numerical solutions for the value of a vaccine using the benchmark parameters from Section 4 but varying the elasticity of intertemporal substitution (EIS). The upper two panels show the full-information solution, with the upper right case corresponding to the benchmark $\psi = 1.5$, whereas the left panel lower the EIS to $\psi = 0.15$. There is almost no difference between the two solutions (which verifies the robustness of the conclusions in Section 4 on this dimension). The bottom two panels show the results under partial information. Specifically, results are computed under the assumption that agents' standard deviation of beliefs about the two parameters are equal to their mean beliefs. Comparing the right-hand panels, we see that this degree of parameters uncertainty has the effect of raising the level of wealth agents in the economy would be willing to surrender for a cure in the baseline case of a high EIS by between 7 and 15 percentage points, or up to a factor of three times the full information value. The left hand panels show the same effect, but amplified to an extreme level. With a low intertemporal elasticity, the representative agent would be willing to sacrifice on the order of 50 to 60 percent of accumulated wealth.

An additional computation that our framework can address is the value of a permanent cure.

³⁶Knowing the solution for higher M enables direct evaluation of the jump-terms in (40)-(41). Knowing the solution on the inner edges enables explicit approximation of the first partial derivatives.

Table 7 shows the fraction of wealth agents in the economy would exchange to live in a world with no pandemics. (Formally, this is equivalent to letting λ go to infinity.) As expected, the values now show the same pattern as in Table 6, but exaggerated still further. In this case, eliminating the threat *and* resolving the parameter uncertainty can lead to valuation of 25 to 50 percent for high EIS agents and 60 to 80 percent for low EIS agents.

The latter finding may be counterintuitive based on the common understanding of Epstein-Zin preferences under which agents with $\psi < 1/\gamma$ can be viewed as having a preference for “later resolution of uncertainty.” In the current model, agents facing a pandemic are much worse off with parameter uncertainty. This is verified in Table 8 where we compute the value that agents would pay to resolve parameter uncertainty *without* ending the on-going pandemic.

For both values of the EIS the numbers are again extremely high, and for the low EIS case they are even higher than in the previous table. Apparently, in this economy, low-EIS agents would pay dearly for early resolution of uncertainty. The source of the extreme welfare loss in this case is the endogenous consumption response. Recall that low-EIS agents cut their consumption during a pandemic. With parameter uncertainty this response becomes extreme because agents cannot rule out the worst case scenario that $\lambda \sim 0$, i.e., that there will never be a cure and the pandemic effectively lasts forever. This possibility leads to extreme savings and, consequently, very little utility flow from consumption.

Even with high EIS however, the effect of parameter uncertainty is economically large, and is again due to agents being unable to rule out worst-case scenarios. From a policy perspective, the implication of this finding is that, while working to end the current pandemic is enormously valuable, equally and perhaps even more valuable is anything that resolves uncertainty about the frequency and, especially, the duration of current and future pandemics. In addition to developing cures and vaccines, understanding the fundamental science behind the fight against viral pathogens and investing in the infrastructure for future responses can provide crucial gains to welfare.

6 Conclusion

In this paper, we estimated the value of a “cure” – vaccine for a pandemic – using the joint behavior of stock prices and a novel vaccine progress indicator based on the chronology of stage-by-stage

progress of individual vaccine candidates and related news. We developed a general equilibrium regime-switching model of repeated pandemics and stages of vaccine progress, wherein the representative agent withdraws labor and alters consumption endogenously to mitigate the economic consequences of health risk arising from pandemics. In the resulting asset-pricing framework, we showed that the covariance of stock prices with the vaccine progress indicator gives an indirect estimate of labor contraction during the pandemic relative to the non-pandemic states; in turn, the empirical estimate of the covariance helps pin down the labor contraction which is an important statistic for the value attached by the representative agent to finding a cure.

With standard preferences parameters, the value of a cure turns out to be worth 5-15% of wealth (formally, capital stock in our model). The value of the cure rises sharply when there is uncertainty about the frequency and duration of pandemics. Indeed, we find that the representative agent would be willing to pay as much for resolution of this parameter uncertainty as for the cure absent such uncertainty, an effect that is stronger – not weaker – when agents have a preference for later resolution of uncertainty. An important policy implication is that understanding the fundamental biological and social determinants of future pandemics, for instance, whether pandemics are related to zoonotic diseases triggered more frequently by climate change, may be as important to mitigating their economic impact as resolving the immediate pandemic-induced crisis.

An interesting extension of our regime-switching framework could be one where as the pandemic evolves through various stages of vaccine progression, it may be simultaneously evolving in its own characteristics. For instance, the arrival intensity (ζ) of the health shock might decline due to “herd immunity” building up or its impact on capital stock (Δ) be mitigated due to learning-by-doing in working from home. Such variations across pandemic states would also generate the realistic implication that labor contraction across pandemic states reduces as the pandemic gets “weaker.” Theoretically, this would add richness to the framework we have proposed; empirically, it would require substantially greater statistical power to estimate state-by-state covariance of stock returns with changes in the vaccine progress indicator as there is progression across the pandemic states.

Our empirical work could be extended in several other directions. First, long-short or “factor mimicking” portfolios can be constructed to map into changes in the vaccine progress indicator for use in future asset-pricing tests. Secondly, changes in the vaccine progress indicator may also

be relevant for fixed income markets and expectations of future interest rates; more generally, progress in finding a cure could affect expectations of monetary and fiscal policies, which we did not consider in this paper. Thirdly, we can numerically consider the possibility of setback or regression in progress of the vaccine by allowing λ_d in our state-transition matrix to be greater than zero, a feature that can have significant implications for asset price volatility, and in turn, for options markets. Finally, vaccines may be more readily available for early deployment in some countries (developed ones, for example) versus others; this would imply patterns in sensitivity of cross-country returns to the vaccine progress indicator, which can be teased out in data.

One caveat to our estimate of the value of a cure is that it is essentially a comparative static exercise. In particular, the economy in our model does not possess the technology to actually effect the transition out of a pandemic. In reality, there is a "real option" to invest in vaccine technologies that affects the probability of switching out of a pandemic. It is an interesting open question for future research to embed the vaccine production technologies into the model, allowing policy analysis that can help answer questions such as: How much should the central planner invest or co-fund the investment in vaccines given their value to the society may far exceed the value to individual vaccine production companies? Should the central planner cap user fees for deployment of the vaccine once developed? How do these choices affect competition in the speed of vaccine development and the endogenous probability of switching out of pandemics? Our asset-pricing perspective on the value of a cure is hopefully a useful first step for further inquiry along these lines.

References


- Asger Lau Andersen, Emilt Toft Hansen, Niels Johannsesen, and Adam Sheridan. Consumer responses to the covid-19 crisis: Evidence from bank account transaction data. Technical report, CEPR, April 2020. URL <https://cepr.org/sites/default/files/news/CovidEconomics7.pdf>.
- 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. URL <http://www.nber.org/papers/w27617>.
- 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. URL <https://www.nber.org/papers/w26949>.
- 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. ISSN 2045-9920. doi: 10.1093/rapstu/raaa008. URL <https://doi.org/10.1093/rapstu/raaa008>. raaa008.
- Robert J. Barro. Rare Disasters and Asset Markets in the Twentieth Century*. *The Quarterly Journal of Economics*, 121(3):823–866, 08 2006. ISSN 0033-5533. doi: 10.1162/qjec.121.3.823. URL <https://doi.org/10.1162/qjec.121.3.823>.
- Robert J Barro. Rare disasters, asset prices, and welfare costs. *American Economic Review*, 99(1): 243–64, 2009.
- Louis-Philippe Beland, Abel Brodeur, and Taylor Wright. Covid-19, stay-at-home orders and employment: Evidence from cps data. Working Paper 13282, IZA Institute of Labor Economics, May 2020. URL <http://ftp.iza.org/dp13282.pdf>.
- Shai Bernstein, Richard R Townsend, and Ting Xu. Flight to safety: How economic downturns affect talent flows to startups. Working Paper 27907, National Bureau of Economic Research, October 2020. URL <http://www.nber.org/papers/w27907>.

- 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. URL <https://ssrn.com/abstract=3568574>.
- 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. URL <http://www.nber.org/papers/w27431>.
- 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 2020a. URL <http://www.nber.org/papers/w27141>.
- Olivier Coibion, Yuriy Gorodnichenko, and Michael Weber. Labor markets during the covid-19 crisis: A preliminary view. Working Paper 27017, National Bureau of Economic Research, April 2020b. URL <http://www.nber.org/papers/w27017>.
- 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. doi: 10.1257/aer.20130392. URL <https://www.aeaweb.org/articles?id=10.1257/aer.20130392>.
- Jonathan I Dingel and Brent Neiman. How many jobs can be done at home? Working Paper 26948, National Bureau of Economic Research, April 2020. URL <http://www.nber.org/papers/w26948>.
- Darrell Duffie and Larry G. Epstein. Asset pricing with stochastic differential utility. 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.

- Eliza Forsythe, Lisa B Kahn, Fabian Lange, and David G Wiczer. Labor demand in the time of covid-19: Evidence from vacancy postings and ui claims. Working Paper 27061, National Bureau of Economic Research, April 2020. URL <http://www.nber.org/papers/w27061>.
- X. Gabaix. Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *Quarterly Journal of Economics*, 127(2):645–700, 2012.
- Harrison Hong, Jeffrey Kubik, Neng Wang, Xiao Xu, and Jinqiang Yang. Pandemics, vaccines and corporate earnings. Technical report, Working Paper, 2020a.
- Harrison Hong, Neng Wang, and Jinqiang Yang. Implications of stochastic transmission rates for managing pandemic risks. Working Paper 27218, National Bureau of Economic Research, May 2020b. URL <http://www.nber.org/papers/w27218>.
- Raymond Kan and Cesare Robotti. On moments of folded and truncated multivariate normal distributions. *Journal of Computational and Graphical Statistics*, 26(4):930–934, 2017.
- Leonid Kogan, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. Technological innovation, resource allocation, and growth. *The Quarterly Journal of Economics*, 132(2):665–712, 2017.
- Julian Kozlowski, Laura Veldkamp, and Venky Venkateswaran. Scarring body and mind: The long-term belief-scarring effects of covid-19. Technical report, 2020 Jackson Hole Economic Policy Symposium Proceedings, 2020.
- Robert Lucas. *Models of Business Cycles*. Blackwell, Oxford, 1987.
- Simon Mongey, Laura Pilossoph, and Alex Weinberg. Which workers bear the burden of social distancing policies? Working Paper 27085, National Bureau of Economic Research, May 2020. URL <http://www.nber.org/papers/w27085>.
- John Muellbauer. The coronavirus pandemic and us consumption. *Vox EU*, 2020. URL <https://voxeu.org/article/coronavirus-pandemic-and-us-consumption>.
- A. O'hagan. Bayes estimation of a convex quadratic. *Biometrika*, 60(3):565–571, 1973.


- Robert S. Pindyck and Neng Wang. The economic and policy consequences of catastrophes. *American Economic Journal: Economic Policy*, 5(4):306–39, November 2013. doi: 10.1257/pol.5.4.306. URL <https://www.aeaweb.org/articles?id=10.1257/pol.5.4.306>.
- 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. ISSN 0027-8424. doi: 10.1073/pnas.2010068117. URL <https://www.pnas.org/content/117/34/20468>.
- 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. doi: 10.1146/annurev-financial-111914-041906. URL <https://doi.org/10.1146/annurev-financial-111914-041906>.
- 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. ISSN 1465-4644. doi: 10.1093/biostatistics/kxx069. URL <https://doi.org/10.1093/biostatistics/kxx069>.

Figure 1: Expected Time to Vaccine Deployment



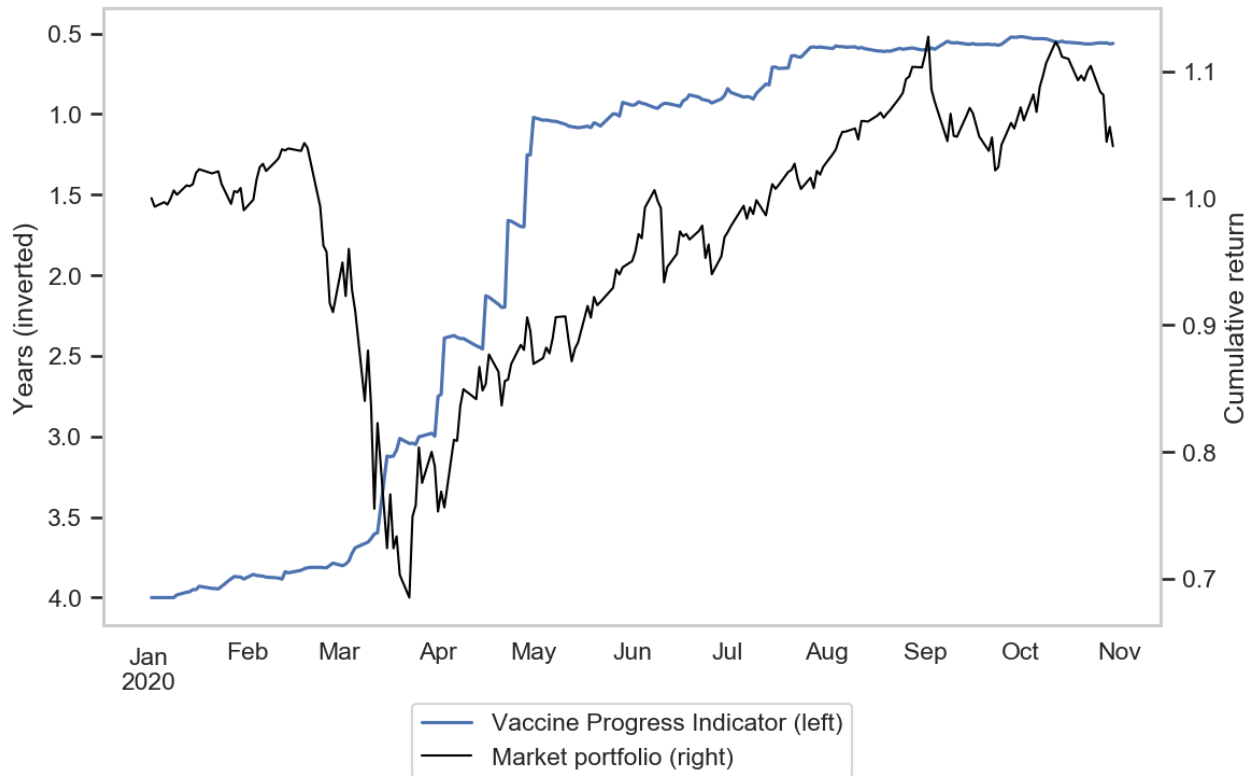
Note: Figure shows our estimate of the expected time to widespread deployment of a COVID-19 vaccine in years. Dashed lines show one standard deviation bands.

Figure 2: Number of Active COVID-19 Vaccine Projects



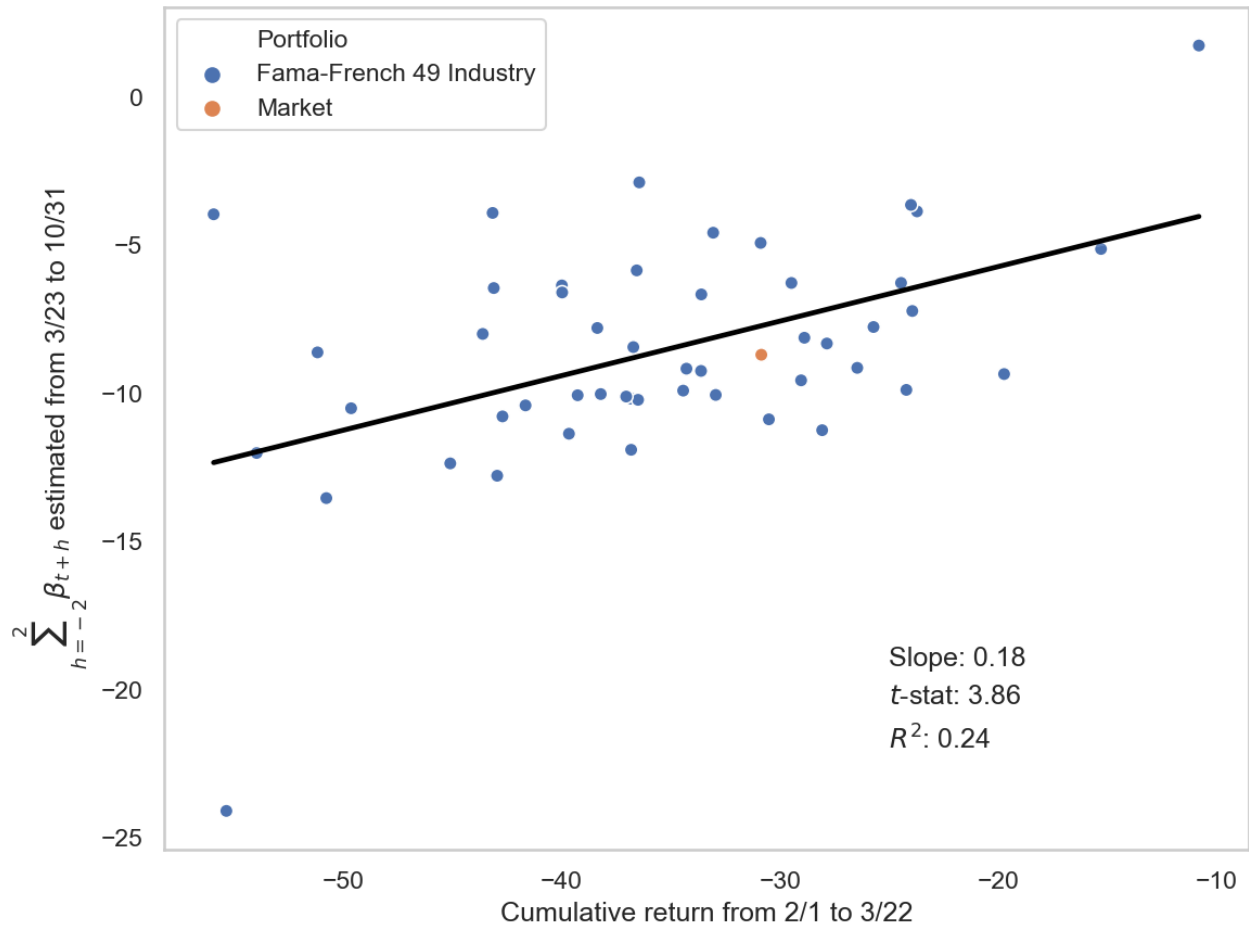
Note: Figure shows the number of active COVID-19 vaccine candidates. Data as of November 2020.

Figure 3: Vaccine Progress and Market Performance




Note: Figure plots vaccine progress (inverted and left axis) along with the cumulated year-to-date excess return on the value-weight CRSP index (right axis). The risk-free rate is the one-month Treasury bill rate.

Figure 4: 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 (3).

Figure 5: Marginal Propensity to Consume in Pandemic and Non-Pandemic States




Note: Figure illustrates for a 10-state regime-switching model that $c(s)$ in the pandemic states, $s \in \{1, \dots, 9\}$, is below (above) that in the non-pandemic state, $s = 0$ or $s = 10$, when ψ is below (above) 1. Parameters chosen are in Table 4.

Figure 6: Stock Market Sensitivity to Vaccine Progress



Note: Figure shows the sensitivity $-\Delta \log P / \Delta E[T^*]$ as a function of g_1 for model solutions varying the pandemic intensity parameters η and λ and the risk neutral expected growth rate of output as described in the text. Each star corresponds to a single model solution.

Figure 7: Value of a Cure



Note: Figure shows the value V as a function of the pandemic intensity parameters η and λ . The left panel plots V against $1/\lambda$. The right panel plots V against η . The left panel sets $\eta = 0.02$ and the right panel sets $\lambda = 0.5$. Both panels take the current state as $s = 1$.

Table 1: Forecast Comparison

<i>Deutsche Bank</i>				
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	
<i>Superforecasters</i>				
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: 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 2: Stock Market Sensitivity to Vaccine Progress News

	(1) OLS	(2) KPSS (Prior 1)	(3) KPSS (Prior 2)
γ_1	-0.070 (0.067)	-0.088 (0.035)	-0.096 (0.035)
γ_2	0.131 (0.092)	0.163 (0.035)	0.168 (0.035)
β_{t-2}	1.316 (1.526)	-0.536 (0.423)	-0.382 (0.290)
β_{t-1}	-4.124 (3.121)	-1.924 (0.746)	-1.269 (0.586)
β_t	-1.100 (0.739)	-0.942 (0.592)	-0.991 (0.606)
β_{t+1}	0.719 (2.054)	-0.517 (0.412)	-0.432 (0.342)
β_{t+2}	-5.404 (1.729)	-2.446 (0.760)	-1.011 (0.458)
α	0.204 (0.097)	0.240 (0.079)	0.279 (0.078)
$\sum_{h=-2}^2 \beta_{t+h}$	-8.593 (0.653)	-6.365 (1.345)	-4.086 (1.056)
N	206	206	206

Note: Table shows the results from regression (2). The dependent variable is daily excess returns on the market portfolio in percent. 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 return on the value-weighted CRSP index is used from January 1, 2020 to October 31, 2020. All columns are employ 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 1 estimates the regression using OLS. Columns 2 and 3 employ the methodology of Kogan et al. (2017) and assume the pre-truncated normal distribution for β_t has standard deviation equal to 1. Column 2 further uses the same prior for all response coefficients, while column 3 uses a pre-truncated standard deviation of 0.7 for the first lead and lag and 0.5 for the second lead and lag. OLS results display Newey-West standard errors with four lags in parentheses and standard deviation of the F -statistic on $\sum_{h=-2}^2 \beta_{t+h}$. KPSS results show posterior standard deviations in parentheses.

Table 3: Stock Market Sensitivity to Vaccine Progress News – 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
$\pi_{\text{approval}}^{\text{base}}$	0.95	0.95	0.95	0.95	0.85	0.95
Ex-China and Russia	Y	Y	Y	Y	Y	N
γ_1	-0.070 (-1.04)	-0.067 (-1.01)	-0.068 (-1.02)	-0.075 (-1.10)	-0.073 (-1.08)	-0.080 (-1.50)
γ_2	0.131 (1.43)	0.116 (1.32)	0.127 (1.42)	0.131 (1.42)	0.134 (1.46)	0.111 (1.37)
β_{t-2}	1.316 (0.86)	2.389 (1.07)	1.543 (0.90)	1.275 (0.73)	0.959 (0.63)	1.980 (1.19)
β_{t-1}	-4.124 (-1.32)	-5.400 (-1.36)	-3.168 (-1.21)	-3.566 (-1.16)	-3.927 (-1.31)	-5.331* (-1.80)
β_t	-1.100 (-1.49)	-0.570 (-0.50)	-1.046 (-1.41)	-1.185 (-1.57)	-1.157 (-1.56)	1.084 (0.78)
β_{t+1}	0.719 (0.35)	2.085 (0.79)	1.112 (0.59)	0.807 (0.43)	0.600 (0.30)	-0.696 (-0.44)
β_{t+2}	-5.404*** (-3.13)	-7.310*** (-4.30)	-5.189*** (-3.65)	-4.872*** (-2.84)	-5.057*** (-2.72)	-4.171 (-1.61)
α	0.204** (2.11)	0.195* (1.94)	0.226** (2.28)	0.220** (2.27)	0.203** (2.11)	0.210** (2.14)
Jump dummies	Y	Y	Y	Y	Y	Y
$\sum_{h=-2}^2 \beta_{t+h}$	-8.593	-8.806	-6.746	-7.541	-8.582	-7.134
F-stat	8.21	5.50	5.37	5.52	8.62	3.69
P-value	0.00	0.02	0.02	0.02	0.00	0.06
N	206	206	206	206	206	206

Note: Table shows the results from regression (2). The dependent variable is daily excess returns on the market portfolio in percent. 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; column 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; column 4 doubles the base copula correlation to 0.4; column 5 decreases the probability of success to 0.85 in the application state; and column 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; t -statistics are shown in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Parameter Values

Parameter	Symbol	Value
Coefficient of relative risk aversion	γ	4.0
Elasticity of intertemporal substitution	ψ	1.5
Rate of time preference	ρ	0.04
Non-pandemic expected output growth	μ	0.055
Non-pandemic output volatility	σ	0.05

Note: Table shows parameter values used in estimating the value of a cure.

Table 5: Value of a Cure: The Effect of Externality

		<u>Central Planner</u>			<u>Benchmark</u>			
		λ			λ			
		0.2	0.5	1.0	0.2	0.5	1.0	
	0.01	0.197	0.094	0.048	0.01	0.242	0.116	0.058
η	0.05	0.154	0.084	0.045	η	0.05	0.185	0.102

Note: Table shows the fraction of wealth that the representative would be willing to surrender for a one-time transition out of the pandemic state. The right panel shows the results when the labor supply decision is made by individual agents acting atomistically. The left panel shows the case where the labor policy is determined by a central planner. All cases use $\gamma = 4, \psi = 1.5, \rho = 0.04, \alpha = 0.5, \sigma = 0.05, \mu = 0.055, \Delta = 0.06, \epsilon = 0.4, k = 0.1, K = 0.4$ and $\zeta = 1$.

Table 6: Value of a Cure under Parameter Uncertainty

		Low Uncertainty / Low EIS			Low Uncertainty / High EIS			
		$\hat{\lambda}$			$\hat{\lambda}$			
		0.2	0.5	1.0	0.2	0.5	1.0	
$\hat{\eta}$	0.01	0.242	0.114	0.058	0.01	0.242	0.116	0.058
	0.05	0.192	0.102	0.055	0.05	0.185	0.102	0.055
		High Uncertainty / Low EIS			High Uncertainty / High EIS			
		$\hat{\lambda}$			$\hat{\lambda}$			
		0.2	0.5	1.0	0.2	0.5	1.0	
$\hat{\eta}$	0.01	0.633	0.613	0.558	0.01	0.379	0.302	0.222
	0.05	0.456	0.479	0.477	0.05	0.256	0.222	0.186

Note: Table shows the fraction of wealth that the representative agent would be willing to surrender for a one-time transition out of the pandemic state. The cases labeled High EIS set $\psi = 1.5$. Cases labeled Low EIS set $\psi = 0.15$. Cases labeled Low Uncertainty correspond to agents knowing the parameters λ and η . Cases labeled High Uncertainty correspond to agents having a posterior standard deviation for those parameters that is equal to their point estimates of them. All cases use $\gamma = 4, \rho = 0.04, \alpha = 0.5, \sigma = 0.05, \mu = 0.05, \Delta = 0.06, \epsilon = 0.4, k = 0.1, K = 0.4$ and $\zeta = 1$.

Table 7: Value of a Permanent Cure

		Low Uncertainty / Low EIS			Low Uncertainty / High EIS			
		$\hat{\lambda}$			$\hat{\lambda}$			
		0.2	0.5	1.0	0.2	0.5	1.0	
$\hat{\eta}$	0.01	0.308	0.136	0.068	0.01	0.327	0.148	0.074
	0.05	0.430	0.214	0.111	0.05	0.429	0.239	0.130
		High Uncertainty / Low EIS			High Uncertainty / High EIS			
		$\hat{\lambda}$			$\hat{\lambda}$			
		0.2	0.5	1.0	0.2	0.5	1.0	
$\hat{\eta}$	0.01	0.813	0.720	0.613	0.01	0.503	0.378	0.265
	0.05	0.831	0.751	0.658	0.05	0.538	0.435	0.335

Note: Table shows the fraction of wealth that the representative agent would exchange to live in a world with no pandemics. High uncertainty denotes agents having a posterior standard deviation for the regime parameters λ and η that is equal to their point estimates of them. Low uncertainty denotes full information. The cases labeled High EIS set $\psi = 1.5$. Cases labeled Low EIS set $\psi = 0.15$. All cases use $\gamma = 4, \rho = 0.04, \alpha = 0.5, \sigma = 0.05, \mu = 0.05, \Delta = 0.06, \epsilon = 0.4, k = 0.1, K = 0.4$ and $\zeta = 1$.

Table 8: Value of Information

		Low EIS			High EIS				
		$\hat{\lambda}$			$\hat{\lambda}$				
		0.2	0.5	1.0	0.2	0.5	1.0		
$\hat{\eta}$	0.01	0.733	0.675	0.587	$\hat{\eta}$	0.01	0.270	0.273	0.209
	0.05	0.708	0.682	0.617		0.05	0.200	0.255	0.236

Note: Table shows the fraction of wealth that the representative would be willing to surrender for a one-time transition from high parameter uncertainty to low parameter uncertainty. High uncertainty denotes agents having a posterior standard deviation for the regime parameters λ and η that is equal to their point estimates of them. Low uncertainty denotes full information. The cases labeled High EIS set $\psi = 1.5$. Cases labeled Low EIS set $\psi = 0.15$. All cases use $\gamma = 4, \rho = 0.04, \alpha = 0.5, \sigma = 0.05, \mu = 0.05, \Delta = 0.06, \epsilon = 0.4, k = 0.1, K = 0.4$ and $\zeta = 1$.

Appendix

Appendix [A](#) includes news articles from the Introduction and Section [3](#).

Appendix [B](#) describes the simulation, data and parameters for the vaccine progress indicator from Section [3](#).

Appendix [C](#) contains proofs to Section [4](#).

An [online appendix](#) includes additional details on the vaccine progress indicator from Section [3](#). It further derives the solution to the regime-switching model with just two states, Proposition [5](#) in Section [4](#), and [\(40\)-\(41\)](#) in Section [5](#).

A News Articles

This section includes news articles from the Introduction and Section [3](#).

A.1 News Articles from the Introduction

On May 18, 2020 *Moderna* released positive interim clinical data from their Phase I trials and announced a Phase III trial.

Federal Reserve chair Jay Powell has warned that a full US economic recovery may take until the end of next year and require the development of a COVID-19 vaccine: "For the economy to fully recover, people will have to be fully confident. And that may have to await the arrival of a vaccine", Mr. Powell told CBS News on Sunday.

[Lauren Fedor and James Politi, Financial Times, May 18, 2020](#)

U.S. stocks gained about \$1 trillion of market capitalization yesterday, and while there are lots of reasons why any particular stock may have gone up or down, good news about a vaccine that might allow reopening of the economy seems like a common factor for a lot of stocks.

"U.S. Stocks Surge as Hopes for Coronavirus Vaccine Build," was the Wall Street Journal's headline, citing the Moderna results... It is almost fair to say that Moderna added \$1 trillion of value to all the other stocks yesterday.

[Matt Levine, Money Stuff, May 19, 2020](#)

On July 14, 2020 *Moderna* publishes positive Phase I data in the *New England Journal of Medicine*, highlighted by its vaccine candidate producing antibodies in all patients.

The most interesting correlation in the stock market right now is the one between (1) the prices of airline stocks and (2) the amount of antibodies produced by coronavirus vaccine candidates in clinical trials. So far the vaccines are experimental and uncertain. If you knew that they'd work really well—protect everyone perfectly, no side effects, easy to produce, etc.—then you'd know with a pretty high degree of certainty that airline stocks (and cruise ships, hotels, casinos, retailers, etc.) would go up. If you knew that they'd be a disaster then you'd probably be short airlines.

So on Tuesday Moderna announced good news, and yesterday... Royal Caribbean Cruises Ltd. was up 21.2%. Norwegian Cruise Line Holdings Ltd. was up 20.7%. Carnival Corp. was up 16.2%. American Airlines Group Inc. was also up 16.2%. United Airlines Holdings was up 14.6%. The biggest gainers were the vaccine sensitive industries, not Moderna itself.

[Matt Levine, Money Stuff, July 16, 2020](#)

On November 9, 2020 *Pfizer* and *BioNTech* announced positive news regarding interim analysis from their Phase III Study.

Markets received a shot in the arm Monday from Pfizer Inc. and its encouraging Stage III tests on a COVID-19 vaccine. As a result, the S&P 500, the MSCI World and the MSCI All-World indexes all rose to records. But that misses the point of the impact. The news triggered the biggest single-day market rotation I've witnessed in the 30 years since I started covering markets...

In technical terms, the clearest expression of the violence of the turnaround comes from tracking the performance of stocks that have had the greatest positive momentum, relative to the market. Bloomberg's measure of the pure momentum factor in the U.S. stock market shows that momentum dropped 4% Monday. Since Bloomberg started tracking daily moves in 2008, it had never before fallen as much as 2%.

[John Authers, Bloomberg Opinion, November 10, 2020](#)

Monday’s news that a COVID-19 vaccine being developed by Pfizer and Germany’s BioNTech was more than 90 per cent effective sent markets soaring. But it also prompted an abrupt switch out of sectors that have prospered during the pandemic, such as technology, and into beaten-down stocks such as real estate and airlines — and triggered an earthquake in some popular investment “factors” such as value and momentum...

The value factor, which is centered on lowly-priced, unfashionable stocks, enjoyed a 6.4 per cent uplift, its strongest one-day gain since the 1980s, while the momentum factor — essentially stocks on a hot streak — tumbled 13.7 per cent, its worst ever loss, according to JPMorgan.

[Laurence Fletcher and Robin Wigglesworth, Financial Times, November 14, 2020](#)

A.2 News Articles from Section 3

Our duration estimates are based on projections from the pharmaceutical and financial press during 2020. For example, see (1) [Damian Garde, STAT News, January 24, 2020](#), (2) [Chelsea Weidman Burke, BioSpace, February 17, 2020](#), (3) [Hannah Kuchler, Clive Cookson and Sarah Neville, Financial Times, March 5, 2020](#), (4) [Bill Bostock, Business Insider, April 1, 2020](#), (5) [Derek Lowe, Science Translational Medicine, April 15, 2020](#), (6) [The Economist, April 16, 2020](#), (7) [Nicoletta Lanese, Live Science, April 16, 2020](#), and (8) [James Paton, Bloomberg, April 27, 2020](#).

B Vaccine Progress Indicator

This section describes the simulation, data and parameters for the vaccine progress indicator. More details are included in the online appendix.

B.1 Simulation Procedure

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

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

and each state has known expected duration τ_s and baseline probability of success π_s^{base} .

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 $\Delta\pi : \rightarrow [-1,1]$ be a mapping from news to changes in probabilities. For each candidate, we cumulate the changes in probabilities from all news from the beginning of our sample t_0 up to time t ,

$$\Delta\pi_{n,t}^{\text{news}} = \sum_{t'=t_0}^t \Delta\pi(\omega_{n,t'}). \quad (\text{A.1})$$

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

$$\pi_{n,s,t}^{\text{total}} = \frac{\exp Y_{n,s,t}}{1 + \exp Y_{n,s,t}} \quad (\text{A.2})$$

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

Figure A.1 outlines the simulation procedure. We simulate stage-by-stage progress of each candidate and generate the expected time to first vaccine deployment, similar to a first to "default" model. Specifically, on each day, one run of the simulation repeats steps one to three until candidates have all failed or deployed:

1. Draw two N -dimensional multivariate Normal random variables

$$z_t^u, z_t^d \sim \mathcal{N}(0, \mathcal{R}) \quad (\text{A.3})$$

2. For each candidate, transform to exponentially driven time to success and failure,

$$t_{n,s,t}^u = -\frac{\log \Phi(z_{n,t}^u)}{\lambda_{n,s,t}^u} \quad \text{and} \quad t_{n,s,t}^d = -\frac{\log \Phi(z_{n,t}^d)}{\lambda_{n,s,t}^d} \quad (\text{A.4})$$

where

$$\lambda_{n,s,t}^u = \frac{\pi_{n,s,t}^{\text{total}}}{\tau_s} \quad \text{and} \quad \lambda_{n,s,t}^d = \frac{1 - \pi_{n,s,t}^{\text{total}}}{\tau_s} \quad (\text{A.5})$$

3. If $t_{n,s,t}^u > t_{n,s,t}^d \implies$ candidate's run is over. Else, candidate advances states, continue run

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

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

5. Then calculate minimum time to vaccine deployment across candidates as $T_m^* = \min_n T_n$

That finishes one run of the simulation. Repeat for $M = 50,000$ runs and then advance to $t + 1$.

On each day across runs, we calculate the average

$$\mathbb{E}[T^*] = (1 - \mu)T_t^D + \mu T^{ND}, \quad (\text{A.6})$$

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.

B.2 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.¹ We observe the start dates of each pre-clinical and clinical trial, along with their vaccine strategy. Vaccines typically take years to 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.1 lists our parameter choices of durations and baseline probabilities of success. Table A.2 summarizes the distribution of time spent in each state in our simulation. Following Wong

¹This version of the paper uses the timeline available on November 2, 2020.

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.3 lists the news types along with their changes in probabilities.

C Proofs to Section 4

C.1 Proof of Proposition 1

Proof. From the evolution of capital stock for the representative agent (16), we obtain the Hamilton-Jacobi-Bellman (HJB) equation as follows for each state $s \in \{1, \dots, S-1\}$

$$0 = \max_{C,l} \left[f(C, \mathbb{J}(s)) - \rho \mathbb{J}(s) + \mathbb{J}_q(s)(l^\alpha q \mu - C) + \frac{1}{2} \mathbb{J}_{qq}(s) l^\alpha q^2 \sigma^2 + \xi [\mathbb{J}(s)(q(1 - \chi \Delta)) - \mathbb{J}(s)(q)] \right. \\ \left. + \lambda_u(s) [\mathbb{J}(s+1)(q) - \mathbb{J}(s)(q)] + \lambda_d(s) [\mathbb{J}(s-1)(q) - \mathbb{J}(s)(q)] \right] \quad (\text{A.7})$$

Using the conjecture for the objective function (17) 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

$$\mathbb{J}_q(q) \alpha l^{\alpha-1} \mu q + \frac{1}{2} \mathbb{J}_{qq}(q) \alpha l^{\alpha-1} \sigma^2 q^2 - \mathbb{J}_q(q(1 - \chi \Delta)) \xi \varepsilon \Delta q = 0 \quad (\text{A.8})$$

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

$$\left[\frac{\alpha \left(\mu - \frac{1}{2} \gamma \sigma^2 \right)}{\xi \varepsilon \Delta} \right] l^{\alpha-1} - [1 - \chi \Delta]^{-\gamma} = 0 \quad (\text{A.9})$$

where $\chi(l, L) = k + \varepsilon l + KL$. In rational expectations equilibrium $L(s) = l(s)$, which gives us that optimal labor in pandemic state $L^*(s) \forall s \in \{1, \dots, S-1\}$ satisfies (20):

$$\chi(L(s), L(s)) = k + (\varepsilon + K)L(s) = \frac{1}{\Delta} \left[1 - (L(s))^{\frac{1-\alpha}{\gamma}} \nu \right] \quad (\text{A.10})$$

where

$$v \equiv \left[\frac{\alpha \left(\mu - \frac{1}{2} \gamma \sigma^2 \right)}{\xi \varepsilon \Delta} \right]^{-1/\gamma}. \quad (\text{A.11})$$

The second-order condition with respect to l is satisfied (footnote 7, equation 19) whenever $(\mu - \frac{1}{2} \gamma \sigma^2) > 0$. For the non-pandemic state $s = 0$ or $s = S$, the third term in first-order condition (A.8) is absent; therefore, we obtain that labor is at the highest possible level $L(0) = L(S) = \bar{l}$, whenever $\alpha (\mu - \frac{1}{2} \gamma \sigma^2) > 0$. \square

C.2 Proof of Propositions 2 and 3

Proof. Taking the first-order condition with respect to $C(s)$ in HJB equation (A.7), we obtain

$$f_c(C, \mathbb{J}(s)) - \mathbb{J}_q(s) = 0. \quad (\text{A.12})$$

Using $f(C, \mathbb{J})$ from (14) and taking the derivative with respect to C , we obtain

$$f_c = \frac{\rho C^{-\psi-1}}{[(1-\gamma)\mathbb{J}(s)]^{\frac{1}{\theta}-1}}, \quad (\text{A.13})$$

which substituting for conjecture $\mathbb{J}(s)$ in equation (16) yields

$$f_c = \frac{\rho C^{-\psi-1}}{H(s)^{\frac{\gamma-\psi-1}{1-\gamma}} q^{\gamma-\psi-1}}. \quad (\text{A.14})$$

Then, for state $s \in \{0, \dots, S\}$, we obtain by substituting $\mathbb{J}_q(s)$ in (A.12), and simplifying:

$$C(s) = \frac{H(s)^{-\theta\psi-1} q}{\rho^{-\psi}}, \quad (\text{A.15})$$

which proves Proposition 3.

To obtain the solution to state-by-state constants $H(s)$, we

1. substitute the optimal controls $\{C(s), L(s)\}$ into the HJB equation (A.7) for each s ;
2. cancel the terms in q which have the same exponent; and

3. group terms not involving $H(s)$ constants into $g(\bar{\ell}, 0)$ for state $s = 0$ and $g(L(s), \xi)$ for state $s \in \{1, \dots, S - 1\}$

to reach equations (22) - (24). 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-pandemic states. \square

The detailed derivation of these equations for the two-state case ($S = 2$) is provided for illustration in the online appendix where we refer to the non-pandemic state as state 0 and the pandemic state as state 1. \square

C.3 Proof of Proposition 4

Proof. The value of a cure (vaccine) $V(s)$ satisfies:

$$\mathbb{J}(0)(q) = \mathbb{J}(0) [(1 - V(s))q] \quad (\text{A.16})$$

where $\mathbb{J}(0)$ is evaluated at $(1 - V(s))q$. Substituting for $\mathbb{J}(s)$ from (17), we obtain

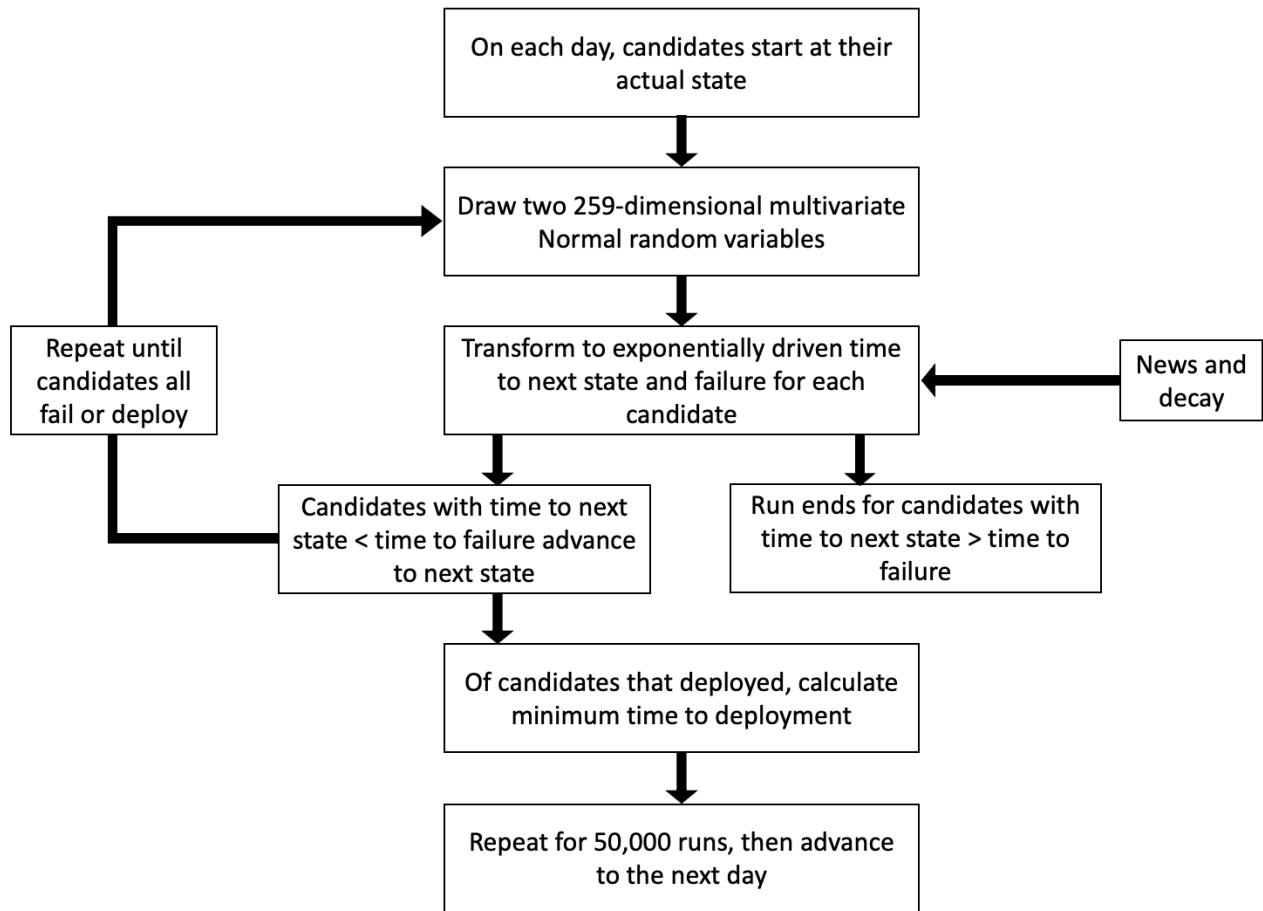
$$\frac{H(0)q^{1-\gamma}}{(1-\gamma)} = \frac{H(0) [(1 - V(s))q]^{1-\gamma}}{(1-\gamma)} \quad (\text{A.17})$$

which yields

$$V(s) = 1 - \left(\frac{H(s)}{H(0)} \right)^{\frac{1}{1-\gamma}}. \quad (\text{A.18})$$

Then, substituting for $C(s)$ from (25) and recognizing marginal propensity to consume, $c(s)$, equals $\frac{dC}{dq} = \frac{C(s)}{q}$, yields Proposition 4. \square

Figure A.1: Simulation Flow Chart



Note: Figure sketches the simulation procedure for estimating the expected time until vaccine deployment.

Table A.1: 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

Note: Table shows the duration and probability of success at each state.

Table A.2: Vaccine States

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

Note: Table shows statistics on the number of days spent in each state before transitioning to the next. 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, among candidates that have successfully transitioned to the next state. Data are from the London School of Hygiene & Tropical Medicine's COVID-19 Tracker. Data are as of November 2, 2020.

Table A.3: News and Changes in Probabilities

<u>Positive</u>		<u>Negative</u>	
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.