

Pandemic Lockdown: The Role of Government Commitment*

Christian Moser[†]

Pierre Yared[‡]

July 30, 2021

Abstract

This paper studies lockdown policy in a dynamic economy without government commitment. Lockdown imposes a cap on labor supply, which improves health prospects at the cost of economic output and consumption. A government would like to commit to the extent of future lockdowns in order to guarantee an economic outlook that supports efficient levels of investment into intermediate inputs. However, such a commitment is not credible, since investments are sunk at the time when the government chooses a lockdown. As a result, lockdown under lack of commitment deviates from the optimal policy. Rules that limit a government's lockdown discretion can improve social welfare, even in the presence of noncontractible information. Quantitatively, lack of commitment causes lockdown to be significantly more severe than is socially optimal. The output and consumption loss due to lack of commitment is greater for higher intermediate input shares, higher discount rates, higher values of life, higher disease transmission rates at and outside of work, and longer vaccine arrival times.

Keywords: Coronavirus, COVID-19, SARS-CoV-2, SIRD Model, Optimal Policy, Pandemic Restrictions, Lockdown, Non-Pharmaceutical Interventions, Rules, Commitment, Flexibility

JEL Classification: E61, H12, I18

*We thank the editor, Loukas Karabarbounis, and three anonymous referees for constructive comments that helped improve the paper. For helpful comments, we also thank Andy Atkeson, Andrés Drenik, Émilien Gouin-Bonenfant, Rick Mishkin, Ben Moll, Trish Mosser, Tommaso Porzio, Jesse Schreger, Steve Zeldes, and seminar participants at Columbia University, the CEPR/IHEID Webinar, and the Barcelona GSE Summer Forum. Miguel Acosta, Rachel Williams, and Entian Zhang provided outstanding research assistance. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System. Any errors are our own.

[†]Columbia University, Federal Reserve Bank of Minneapolis, and CEPR. Email: c.moser@columbia.edu.

[‡]Columbia University and NBER. Email: pyared@columbia.edu.

1 Introduction

The COVID-19 pandemic brought about a great rise in both epidemiological and policy uncertainty.¹ In response to the pandemic, governments across the world implemented lockdown policies to limit the spread of infections. In numerous cases, these policies were first scheduled to end in the near future and then were extended. For instance, New York Governor Andrew Cuomo imposed a statewide stay-at-home order on March 22, 2020, with an initial end date of April 19. This lockdown was later extended, first until April 29 and then until May 15. While several restrictions were further extended on May 15, Cuomo also presented a clear contingency plan with criteria for lifting restrictions in the future.² Elsewhere, the discretion to extend lockdowns was limited by decree. For example, on September 25, 2020, Florida Governor Ron DeSantis announced a lower limit of 50 percent on allowed restaurant capacity, regardless of local restrictions. The stated goal of this lower limit was to reduce future lockdown policy discretion by local governments.³ Similar lockdown extensions, rules for lifting them, and restrictions on future lockdowns were implemented by many other regional and national governments.

As is evident from these examples, lockdown policies create additional uncertainty over and above that posed by epidemiological factors. Such uncertainty affects businesses that need to make forward-looking investments subject to sunk costs. Common examples of sunk costs include airlines maintaining their fleet, hotels deciding how many employees to retain on payroll, and restaurants placing inventory orders ahead of reopening. Because these investments are forward looking, lockdown policies dynamically impact current economic activity through businesses' expectations of their government's plans for reopening.

To formalize these dynamics, in this paper, we study the role of government commitment

¹For example, [Baker et al. \(2020\)](#) find that the onset of the pandemic led to a fourfold increase in their Economic Policy Uncertainty index, which reached its highest value on record. Using text analysis of earnings conference calls, [Hassan et al. \(2020\)](#) track firm-level risks and sentiments due to government-related and other factors. A report by McKinsey & Company concludes that "lockdowns also cause uncertainty to remain high" and that "this uncertainty is paralyzing" ([Smit et al., 2020](#)).

²The New York Forward initiative lays out a detailed guide to reopening businesses, sending people back to work, and allowing social gatherings.

³In an essay published in the Wall Street Journal, DeSantis pleaded the case for policy commitments to preserve government credibility:

Perhaps most damaging to public trust was the public-health campaign urging "15 Days to Slow the Spread." This short-term mitigation, we were told, was necessary to buy time to prepare hospitals for any patient surges. But that reasonable aim was soon transformed into a lockdown-until-eradication approach that left no end in sight for most Americans. Going from "save the hospitals" to "zero Covid" represents one of the greatest instances in history of moving the goal post. ([DeSantis, 2021](#))

in designing lockdown policy. We consider a dynamic economy that embeds sequential government policy decision-making into a general SIRD model of pandemics (Kermack and McKendrick, 1927; Ferguson et al., 2020). Each period, firms invest in intermediate inputs before the government chooses a lockdown policy and workers supply labor. A lockdown imposes an upper bound on labor supply, limiting disease spread at the cost of economic activity. Our framework is general and subsumes key mechanics of many macroeconomic SIRD models in the literature with lockdown or disease-mitigation policies.⁴ A key feature of our model is that investment in intermediate inputs is determined before a lockdown policy is chosen. We think of this as capturing the kinds of investments that businesses make in maintenance, employee retention, and inventory while anticipating the ensuing trajectory of lockdown policies during a pandemic. Through the forward-looking nature of investment, current economic activity depends on firms' expectations of future lockdown policy.

Lockdowns induce both health benefits as well as output and consumption costs. In our model, lockdown reduces contemporaneous disease spread during a pandemic, which evolves according to a modified SIRD model. At the same time, through two channels, output and consumption decrease with the intensity of the lockdown. First, they decrease statically, as labor supply is directly curbed by the lockdown. Second, they decrease dynamically through lower investment in anticipation of lower future marginal returns to investment resulting from future lockdown. Under government commitment, the optimal lockdown policy equates its marginal health benefits with the output and consumption costs.

Our main result concerns the effect of a government's lack of commitment on optimal lockdown policy. A government would like to commit to limit the extent of future lockdowns in order to support more optimistic firm expectations in the present. However, such a commitment is not credible, since investment decisions are sunk when the government decides on future lockdowns. Faced with a sunk investment, a government without commitment wants to impose a more stringent lockdown relative to the optimal policy under commitment, because it does not fully internalize the associated reduction in returns to investment in intermediate inputs. Firms rationally foresee the government's lack of commitment, causing them to invest less than they would in anticipation of the policy under commitment. Through this mechanism, lack of commit-

⁴See Alvarez et al. (2020), Atkeson (2020a,b), Berger et al. (2020), Chari et al. (2020), and Eichenbaum et al. (2020a,b), among others.

ment distorts the efficient levels of investment and therefore output and consumption associated with lockdown policy.

In light of this time inconsistency problem, we study how a government can improve the efficiency of lockdown policy by committing ex-ante to a contingent plan that depends on the evolving health state. We show that an ex-ante rule that imposes state-contingent limits on future lockdown severity can attain the efficient allocation.

We extend the model to a setting in which additional information arrives during a lockdown. Examples of such information include estimates of disease mortality, the state of the economy, the medical system's capacity, or progress on vaccine development. Some of this information may be relevant for the payoffs and costs of lockdown policy. If this information is a contractible part of the state space, we show that it continues to be the case that an ex-ante rule that imposes state-contingent limits on future lockdown severity can attain the efficient allocation. Moreover, even if this information is not contractible—so that policy flexibility is valuable—rules that limit lockdown severity increase social welfare. This is because it is always socially beneficial on the margin to prevent excessive future lockdowns as a means of raising investment in the present.

These results provide a theoretical justification for the social benefits of mandated limits on future lockdowns, such as those implemented by some state governments in the United States. It is important to note that our analysis does not imply that lockdowns are harmful. In fact, reducing or lifting the lockdown is detrimental if the associated health costs exceed the economic gains. However, committing to limiting future lockdowns is beneficial if the economic gains from stimulating investment toward its efficient level exceed the health costs.⁵

In a quantitative exercise, we use a calibrated version of our model to show that lack of commitment leads to an overly severe lockdown, with significant output and consumption losses compared with those of the policy under commitment. We show that the output and consumption losses are greater for higher discount rates, higher values of life, higher disease transmission rates, higher intermediate input shares, and longer vaccine arrival times. Our findings suggest that optimal policy commitments to limit lockdown would result in a significant reduction of out-

⁵Naturally, there are other reasons why a government may choose inefficiently lax lockdowns. Our model abstracts from policy biases involving insufficient degrees of lockdowns by assuming that policies are chosen by a rational and benevolent government that maximizes long-run social welfare. This assumption may be violated if political economy considerations lead the government to overweigh immediate economic gains relative to future health costs of relaxing a lockdown, akin to the mechanism in [Aguiar and Amador \(2011\)](#). The mechanism we highlight in our paper would act against political economy considerations that lead to departures from the assumption of a benevolent government.

put and consumption losses during a pandemic.

Related literature. This paper relates to the nascent literature on optimal policy in a pandemic and, in particular, to the work of [Alvarez et al. \(2020\)](#), [Atkeson \(2020a,b\)](#), [Berger et al. \(2020\)](#), [Chari et al. \(2020\)](#), and [Eichenbaum et al. \(2020a,b\)](#). This literature focuses on the optimal design of government policy, including the timing and intensity of lockdowns, under the assumption that the optimal policy can be enforced at all dates and under all contingencies. Our work highlights that such analyses omit an important aspect of lockdown design—namely, that the optimal policy may be hard to enforce because of issues of time inconsistency. What distinguishes our approach is the focus on the value of government commitment to lockdown policy and the optimal design of rules that limit government discretion.⁶

That prior work on policy responses to a pandemic has ignored issues of time inconsistency is perhaps surprising, given the parallel insights from an older literature that studies government commitment in the context of capital taxation. This body of work includes the important contributions by [Kydland and Prescott \(1980\)](#), [Chari and Kehoe \(1990\)](#), [Klein et al. \(2008\)](#), [Aguiar et al. \(2009\)](#), and [Chari et al. \(2019\)](#). As it does in the previous work on capital taxation, in our model lack of commitment reduces economic activity by distorting investment. Relative to this literature, our work incorporates two new insights that are central to the context of pandemics. First, a lockdown distorts investment not directly via capital taxation but indirectly by lowering the marginal returns to investment through a cap on labor supply. Since lockdown distorts labor, in a way similar to how a labor income tax does, our work more broadly highlights the existence of a time consistency problem that would arise in a model of labor taxation with endogenous labor supply and capital: A government distorting labor ex-post does not internalize the ex-ante effect on decisions by investors. A second difference relative to the capital taxation literature is that investment distortions from lockdown serve not to relax the government budget constraint but instead to improve the future health state. Since this health state is not static but evolves according to an SIRD model, the tradeoff faced by the government is not static but dynamic, and the time inconsistency problem evolves over time.⁷

⁶Complementary to our focus on public commitment, [Chari et al. \(2021\)](#) study the role of private commitment in an island economy with local externalities.

⁷An additional technical complication arises in the present context: The value of a given health state in our model cannot be represented by a univariate, concave function, as in a typical model of optimal fiscal policy. Therefore, the usual methods for comparative statics do not apply here. Instead, we characterize the time inconsistency of optimal

Our analysis of rules for lockdown policy in the presence of noncontractible information relates to a growing literature on commitment versus flexibility in macroeconomics (Athey et al., 2005; Amador et al., 2006; Halac and Yared, 2014, 2018; Moser and Olea de Souza e Silva, 2019). Prior work in this area has focused on rules for either savings or monetary and fiscal policy. Our work adds to this literature and to a growing number of papers on the economics of pandemics—specifically, the theoretical analysis of optimal lockdown policy. Our result that rules can strictly increase social welfare, even if flexibility is valuable, is reminiscent of similar insights in the context of savings or fiscal and monetary policy. However, our results do not directly follow from the methods developed in prior work, which rely on stronger assumptions on the utility function and the information structure than the ones we require in our setting. By extending these insights and applying them to optimal lockdown design, we highlight an overlooked aspect of the debate around lockdown policy during pandemics.

2 Model

We consider a general infinite-horizon model of an economy during a pandemic. Each period has four stages. First, firms make a costly and irreversible investment in intermediate inputs that enhances future productivity (e.g., expenses related to maintenance, personnel, inventory, rent, utilities, overheads, software licenses, and marketing). Second, after the investment is undertaken, the government chooses a lockdown policy, which imposes a cap on labor supply, thereby inhibiting disease spread while reducing economic output and consumption. Third, production takes place, and all proceeds are paid to firms and workers. Fourth and finally, the pandemic evolves according to an SIRD model of disease spread, which depends on the lockdown policy. A key feature of our model is that investment is determined before lockdown policy is chosen. We think of this feature as capturing the fact that business purchases of irreversible inputs must be made in advance of production and in anticipation of future policies. We will explore the implications of this sequencing of investment and lockdown decisions for the optimal policy under commitment compared with that under lack of commitment.

lockdown policy under weak assumptions on the economic environment and the SIRD model of disease dynamics.

2.1 Economic Environment

Periods are indexed by $t = 0, 1, \dots$. The economy is populated by a continuum of workers of unit mass. The distribution of susceptible, infected, recovered, and deceased agents is summarized by the prevailing health state Ω_t , which we discuss in detail below. At every date t , competitive firms make an irreversible investment x_t . The government then chooses a lockdown policy $L_t \in [0, 1]$ representing the fraction of labor supply that is prohibited from working. If $L_t = 0$, then there is no lockdown and all agents can go to work, while if $L_t = 1$, then there is maximal lockdown and no agent is allowed to work. Agents inelastically supply effective labor ℓ_t up to an upper bound of $(1 - L_t)\bar{\ell}(\Omega_t)$, which depends on lockdown policy through the term $(1 - L_t)$ and on the health state through the term $\bar{\ell}(\Omega_t)$.⁸ Anticipating the labor market clearing condition,

$$\ell_t = (1 - L_t)\bar{\ell}(\Omega_t),$$

we can interchangeably refer to labor supply ℓ_t and lockdown policy L_t given some health state Ω_t .

Workers consume their wage income

$$c_t = w_t \ell_t, \tag{1}$$

where c_t is aggregate consumption and w_t is the equilibrium wage. The irreversible investment x_t , combined with labor ℓ_t , generates gross output y_t , according to the following production technology:

$$y_t = f(x_t, \ell_t, \Omega_t), \tag{2}$$

where Ω_t is the health state at date t that is described in detail in the next subsection. The dependence of the production function $f(\cdot)$ on the health state captures the possibility that the pandemic—in addition to making people sick and killing people—decreases output by debilitating the workforce, by changing the share of the labor force working from the office versus from home (Dingel and Neiman, 2020; Mongey et al., 2020) and by inducing protective but productivity-reducing social distancing efforts even in the absence of any lockdown (Farboodi et al., 2020).

⁸This allows for the possibility that, for example, deceased agents cannot work or that infected agents are effectively less productive at work.

We assume that the function $f(\cdot)$ is continuously differentiable, increasing, and globally concave in x_t and ℓ_t , with $\lim_{x_t \rightarrow 0} \partial f(\cdot) / \partial x = \lim_{\ell_t \rightarrow 0} \partial f(\cdot) / \partial \ell = \infty$ and $\lim_{x_t \rightarrow \infty} \partial f(\cdot) / \partial x = \lim_{\ell_t \rightarrow \infty} \partial f(\cdot) / \partial \ell = 0$. From here on, we make the following key assumption:

Assumption 1. *The production function $f(x_t, \ell_t, \Omega_t)$ satisfies*

$$\frac{\partial^2 f(\cdot)}{\partial x_t \partial \ell_t} > 0. \quad (3)$$

Assumption 1 states that investment x_t and labor ℓ_t are q-complements in production. It implies that there are higher marginal returns to investment x_t when labor ℓ_t is greater and vice versa, which is intuitive under our interpretation that x_t is investment that enhances future productivity.

Firm owners maximize profits

$$\pi_t = y_t - rx_t - w_t \ell_t, \quad (4)$$

where $r > 0$ is the exogenously given price of the irreversible investment x_t and the price of gross output is normalized to 1. In a competitive equilibrium, the marginal product of investment satisfies the following firm optimality condition:

$$\frac{\partial f(x_t, \ell_t, \Omega_t)}{\partial x} = r. \quad (5)$$

Equation (5) implies that in a competitive equilibrium in which the optimal investment adjusts to the anticipated level of labor supply,

$$x_t = x^*(L_t, \Omega_t), \quad (6)$$

where the function $x^*(\cdot)$ satisfies $\partial x^*(L_t, \Omega_t) / \partial L < 0$ by Assumption 1. In other words, as a result of the q-complementarity between investment and labor in production, firms invest less in anticipation of a more stringent lockdown.

Labor is competitively supplied so that wages equal the marginal product of labor given by

$$\frac{\partial f(x_t, \ell_t, \Omega_t)}{\partial \ell} = w_t. \quad (7)$$

From equation (7), consumption in (1) can be written as

$$c_t = c^*(x_t, L_t, \Omega_t), \quad (8)$$

where the function $c^*(\cdot)$ is continuously differentiable in x_t and L_t and strictly increasing in x_t by Assumption 1.⁹

2.2 Disease Spread, Lockdown Policy, and Welfare

We model disease spread as following an SIRD model (Kermack and McKendrick, 1927; Ferguson et al., 2020), which we allow to depend on a lockdown policy, as in Atkeson (2020a), Eichenbaum et al. (2020a), and Alvarez et al. (2020). Specifically, the health state of the economy in period t is summarized by $\Omega_t = \{S_t, I_t, R_t, D_t\}$, where $S_t \in [0, 1]$ is the share of susceptible individuals, $I_t \in [0, 1]$ is the share of infected and contagious individuals, $R_t \in [0, 1]$ is the share of recovered individuals, and $D_t \in [0, 1]$ is the share of deceased individuals. It follows that

$$S_t + I_t + R_t + D_t = 1. \quad (9)$$

An SIRD model defines a mapping $\Gamma(\cdot)$ that implies a law of motion of the health state,

$$\Omega_{t+1} = \Gamma(L_t, \Omega_t),$$

which depends on the degree of lockdown at date t .¹⁰ The initial health state Ω_0 is taken as given.¹¹

Social welfare equals the discounted sum of utility streams,

$$\sum_{t=0}^{\infty} \beta^t u(c_t, \Omega_t), \quad (10)$$

where $\beta \in (0, 1)$ is the discount factor, and $u(\cdot)$ is a strictly increasing and strictly concave utility function of consumption c_t and also depends on the health state Ω_t .

⁹We do not require that $c^*(\cdot)$ be globally increasing in L_t , though this will be the case for commonly used production functions such as those in the Cobb-Douglas family.

¹⁰All of our results extend to a setting in which the health state is a function of time or is stochastic, a feature that would capture factors such as the evolving constraints on the medical system and the changing likelihood of vaccine discovery. Our quantitative exercise considers an environment in which a vaccine arrives in finite time.

¹¹We assume that x_0 is endogenous, implying that it is chosen in anticipation of the government's initial lockdown policy. Our main results are robust to assuming x_0 is exogenous.

To simplify the exposition, we assume that the government puts positive weight on only workers' utility. Our main results require that workers and firm owners be distinct and that the government put greater weight on workers. Therefore, the government does not fully internalize the impact of lockdown on intermediate input investment.¹²

Note that utility depends directly on the health state, which may capture the costs of illness and mortality associated with disease spread. Moreover, utility also indirectly depends on disease spread through the level of consumption c_t , since the health state Ω_t directly enters the production function $f(\cdot)$.

Note that our framework is sufficiently general to accommodate considerations such as endogenous social distancing, which would have an effect on utility through $u(\cdot)$, on gross output through $f(\cdot)$, and on disease spread through $\Gamma(\cdot)$, since these are all functions of the health state. From this perspective, the appropriate interpretation of the lockdown policy L_t is that it corresponds to a binding government mandate above and beyond the endogenous social distancing response. This government restriction can be useful for mitigating disease spread if there is an externality associated with endogenous social distancing, where individuals do not internalize the disease cost of their social interactions.¹³

We do not restrict how the health state and lockdown impact gross output, utility, and disease dynamics in the economy, other than by making the following assumption, which we henceforth maintain:

Assumption 2. *The functions $f(x_t, \ell_t, \Omega_t)$, $u(c_t, \Omega_t)$, and $\Gamma(L_t, \Omega_t)$ are continuously differentiable in all elements of Ω_t .*

This technical assumption guarantees that the government's problem is well behaved and that we can rely on first-order conditions (FOCs) in the proofs of our results. Note that these assumptions are satisfied in many recent macroeconomic models with SIRD modules in which disease dynamics respond smoothly to lockdown policies, such as [Alvarez et al. \(2020\)](#) and [Eichenbaum et al. \(2020a\)](#).

¹²Time inconsistency emerges in the present context because the government does not internalize the cost of lockdown on firm owners. If, instead, workers fully owned the firms, then the optimality of the investment decision would imply that the government's ex-ante and ex-post optimal lockdown choices coincide. To see this, note that optimality implies that both ex-ante and ex-post, a marginal increase in lockdown severity associated with lower investment would have zero marginal social net payoff. Thus, there would be no problem of time inconsistency.

¹³See [Farboodi et al. \(2020\)](#) for a discussion of the interaction between endogenous and government-mandated social distancing.

3 Optimal Policy under Commitment

Suppose that the government commits to an optimal lockdown policy sequence $\{L_t^c\}_{t=0}^\infty$ at time 0. This means that the government internalizes the fact that investment optimally adjusts to anticipated labor supply as determined by future lockdown policy. Given firm optimality in (6), this policy sequence induces sequences of optimal labor supply $\{\ell_t^c\}_{t=0}^\infty$ and investment $\{x_t^c\}_{t=0}^\infty$ under government commitment.

After substituting the investment function $x_t = x^*(L_t, \Omega_t)$ from (6) and the consumption function $c_t = c^*(x_t, L_t, \Omega_t)$ from (8) into the social welfare function (10), the government with commitment solves the following sequence problem:

$$\begin{aligned} & \max_{\{L_t\}_{t=0}^\infty} \left\{ \sum_{t=0}^{\infty} \beta^t u(c^*(x^*(L_t, \Omega_t), L_t, \Omega_t), \Omega_t) \right\} & (11) \\ \text{s.t. } & L_t \in [0, 1], \quad \forall t \geq 0, \\ & \Omega_{t+1} = \Gamma(L_t, \Omega_t), \quad \forall t \geq 0, \\ & \Omega_0 \text{ given.} \end{aligned}$$

It is important to note that substituting the optimal firm investment response $x^*(L_t, \Omega_t)$ into the welfare function before deriving the optimal lockdown sequence $\{L_t^c\}_{t=0}^\infty$ that solves the program in (11) means that in all periods, the government with commitment takes into account investment's reaction to its policies. The problem of the government with commitment can be written recursively as

$$V^c(\Omega) = \max_{L \in [0, 1]} \{u(c^*(x^*(L, \Omega), L, \Omega), \Omega) + \beta V^c(\Gamma(L, \Omega))\}, \quad (12)$$

where $V^c(\Omega)$ denotes the value of health state Ω to the government with commitment. The solution to program (12) induces an optimal lockdown policy under commitment as a function of the prevailing health state Ω , denoted by $L^c(\Omega)$. This lockdown policy in turn yields an optimal investment level under commitment that depends only on the health state Ω , denoted by $x^c(\Omega) = x^*(L^c(\Omega), \Omega)$.

Standard arguments, together with Assumption 2, imply that $V^c(\Omega)$ is continuously differentiable in all elements of Ω . This means that the necessary FOC for interior optimal levels of

lockdown under commitment $L^c \in (0, 1)$ is

$$\frac{\partial u(\cdot)}{\partial c} \left[\frac{\partial c^*(\cdot)}{\partial x} \frac{\partial x^*(\cdot)}{\partial L} + \frac{\partial c^*(\cdot)}{\partial L} \right] = -\beta \frac{dV^c(\cdot)}{dL}. \quad (13)$$

In choosing the degree of lockdown, the government weighs two opposing forces, as in [Gourinchas \(2020\)](#) and [Hall et al. \(2020\)](#). On the one hand, it considers the economic costs captured by the left-hand side of (13). The economic costs are twofold. First, conditional on the level of investment, a lockdown has a direct impact on output and consumption by limiting labor supply. Second, a lockdown has an indirect impact on output and consumption by reducing the marginal product of investment, which lowers the optimal investment level. The government's ability to commit gives it the ability to take into account both of these factors and anticipate firms' reaction to the policy.

On the other hand, the government considers the discounted future health benefits in terms of reduced mortality from inhibiting the disease spread, as captured by right-hand side of (13). Differentiating (12), we can write the marginal health benefits of lockdown recursively as

$$\frac{dV^c(\Omega')}{dL} = \frac{dV^c(\Omega')}{d\Omega'} \frac{d\Gamma(L, \Omega)}{dL} \quad (14)$$

$$= \frac{du(c^*(x^c(L', \Omega'), L', \Omega'), \Omega')}{dL} + \beta \frac{dV^c(\Gamma(L', \Omega'))}{dL}, \quad (15)$$

where $\Omega' = \Gamma(L, \Omega)$ denotes next period's health state and L' denotes the level of next period's optimal lockdown. By use of the envelope theorem, the optimal lockdown policy function $L^c(\Gamma(L, \Omega))$ was replaced with the level of next period's optimal lockdown L' on the right-hand side of equation (15). This equation illustrates that present lockdown dynamically impacts all future health states, which in turn impact welfare both through their direct health costs and through their indirect effect on consumption.

4 Optimal Policy under Lack of Commitment

Under lack of commitment, investment is treated as fixed at the time when lockdown policy is decided on. The government at date t chooses an optimal degree of lockdown that depends on sunk investment x_t and the health state Ω_t ; the degree of lockdown is denoted by $L^*(x_t, \Omega_t)$. Firms in turn anticipate the government's policy and decide on the optimal investment level $x^*(L_t, \Omega_t)$,

which depends on the expected lockdown L_t and the health state Ω_t . We consider a Markov perfect equilibrium (MPE) in which investment and lockdown policy can be expressed as functions of only the health state Ω_t —namely $x^n(\Omega_t)$ and $L^n(\Omega_t)$. In any MPE, $x^n(\Omega_t) = x^*(L^n(\Omega_t), \Omega_t)$ and $L^n(\Omega_t) = L^*(x^n(\Omega_t), \Omega_t)$, as the government and firms take each other's reaction functions as given when choosing their actions under the prevailing health state.

The problem of the government without commitment in an MPE can be written recursively as

$$W^n(x, \Omega) = \max_{L \in [0,1]} \{u(c^*(x, L, \Omega), \Omega) + \beta V^n(\Gamma(L, \Omega))\}, \quad (16)$$

$$V^n(\Omega') = u(c^*(x^n(\Omega'), L^n(\Omega'), \Omega'), \Omega') + \beta V^n(\Gamma(L^n(\Omega'), \Omega')), \quad (17)$$

where $W^n(x, \Omega)$ denotes the value to the government given investment x and health state Ω , while $V^n(\Omega')$ denotes the continuation value to the government given next period's health state $\Omega' = \Gamma(L, \Omega)$ in the absence of future government commitment. Note that $W^n(x, \Omega)$ depends on the current period's investment and health state, while $V^n(\Omega')$ depends only on next period's health state. This reflects the fact that next period's MPE investment function $x^n(\Omega')$ is already consistent with the future MPE lockdown policy $L^n(\Omega')$ by the government without commitment and vice versa. Importantly, by not substituting the current period's optimal investment response when solving its problem, the government without commitment treats current investment as sunk when deciding on lockdown policy.

Consider the government's FOC in a differentiable MPE for interior lockdown $L^n \in (0, 1)$ under lack of commitment:

$$\frac{\partial u(\cdot)}{\partial c} \frac{\partial c^*(\cdot)}{\partial L} = -\beta \frac{dV^n(\cdot)}{dL}. \quad (18)$$

Holding all else—including investment and the health state—fixed, the left-hand side of the optimality condition under lack of commitment in (18) is strictly greater than that under commitment in (13). The reason for this is that $\partial x^*(\cdot)/\partial L < 0$ owing to q-complementarity between x and ℓ , which is given by Assumption 1. This captures the fact that compared with a government with commitment, a government without commitment undervalues the economic cost of a lockdown. Specifically, a government without commitment does not take into account that a more stringent lockdown changes ex-ante firm expectations in a way that reduces the level of investment, thereby

reducing future output and consumption.

Turning to the right-hand side of (18), the derivative of the government's continuation value with respect to lockdown is

$$\frac{dV^n(\Omega')}{dL} = \frac{dV^n(\Omega')}{d\Omega'} \frac{d\Gamma(L, \Omega)}{dL} \quad (19)$$

$$= \frac{du(c^*(x^n(L', \Omega'), L', \Omega'), \Omega')}{dL} + \beta \frac{dV^n(\Gamma(L', \Omega'))}{dL} \quad (20)$$

$$+ \left[\frac{du(c^*(x^n(L', \Omega'), L', \Omega'), \Omega')}{dL'} + \beta \frac{dV^n(\Gamma(L', \Omega'))}{dL'} \right] \frac{dL^n(\Omega')}{dL},$$

where $\Omega' = \Gamma(L, \Omega)$ denotes next period's health state as a function of the current lockdown level and health state, L' denotes the level of optimal lockdown under lack of commitment next period, and $L^n(\Omega')$ is next period's MPE lockdown policy under lack of commitment as a function of next period's health state.

The first line on the right-hand side of equation (20) is analogous to that under commitment in (15). It represents the payoff from changing the future health state by changing the lockdown today, holding fixed the optimal future lockdown policy.

The second line on the right-hand side of equation (20) is unique to the case of lack of commitment. It corresponds to the strategic effect of a lockdown today on future policy, since changing the future health state also changes future lockdown incentives. Under commitment, the term analogous to that in brackets in the second line of (20) is identically zero because the government with commitment takes into account firms' reaction to its lockdown choice, as captured by the FOC (13). Under lack of commitment, however, equation (18) and Assumption 1 together imply that the term in brackets is negative.¹⁴

Note that a complexity associated with this general model is that the value of a given health state cannot be represented by a univariate, concave function, as in typical models of optimal fiscal policy. Nevertheless, under the weak conditions spelled out above, we obtain the following result:

Proposition 1 (Time Inconsistency). *Suppose that the optimal policy under commitment $\{L_t^c\}_{t=0}^\infty$ admits an interior solution in some period t . Then, the optimal policy under commitment is time inconsistent.*

Proof. See Appendix A.1. □

¹⁴While we can sign the term in brackets, we cannot sign the overall strategic effect, since the sign of the term $dL^n(\Gamma(L, \Omega))/dL$ is ambiguous because of the nonlinear dynamics of the SIRD model. If, for example, a marginal increase in L causes a large (small) share of the population to become recovered and immune, then the optimal future $L^n(\Gamma(L, \Omega))$ may decrease (increase).

Proposition 1 states that lack of government commitment may result in an inefficient lockdown policy. The idea behind the proof is as follows: If the optimal lockdown policy under lack of commitment was congruous to that under commitment, then the no-commitment government would have no incentive to deviate, because any deviation would be associated with weakly negative change in welfare. But at an interior solution where $L_t \in (0, 1)$ for some t , the optimality condition (18) under no commitment calls for a strictly higher value of L_t than that in condition (13) under commitment. Therefore, the optimal policy is time inconsistent whenever it is interior.¹⁵

The intuition for this result is that absent commitment, the government treats firm investment as fixed and thus undervalues the economic cost of a lockdown, leading to an inefficient choice of lockdown. By anticipating this behavior, firms invest less than they would if the government had commitment. For this reason, the optimal policy under lack of commitment differs from that under commitment.¹⁶

Note that Proposition 1 does not specify whether the optimal lockdown policy under commitment is more or less stringent than that under lack of commitment. This is due to two key differences between the optimal policies with and without commitment. The first difference is a static one: starting from an L_t that is interior under commitment and given a health state Ω_t , investment x_t , and continuation value $V(\Omega_{t+1})$, a government without commitment would choose a strictly higher L_t than a government with commitment. This is because the government without commitment treats investment x_t as sunk when it decides on lockdown policy at time t . The second difference is a dynamic one: Given the difference in policy functions of governments with and without commitment, investment and the health state will evolve differently in a dynamic model under commitment versus under lack of commitment. Without further model restrictions, this makes it challenging to provide a sharp theoretical characterization of the policy path under commitment versus that under no commitment. In the quantitative exercise presented in Section 6, we use a calibrated version of our model to show that at most points in time along the equilibrium path of a simulated pandemic, lockdown under lack of commitment is more severe than that under full commitment.

¹⁵In the event that there are exogenous limits on lockdown policy, as in Acemoglu et al. (2020), an analogous argument applies whenever the policy is interior relative to such exogenous limits.

¹⁶In the case in which Assumption 1 is reversed—i.e., if x and ℓ are q-substitutes in production—then the result in Proposition 1 continues to hold, but the intuition is also reversed.

5 Value of Rules

We have established that the optimal lockdown policy is time inconsistent. Deviations from the policy under commitment occur because a government without commitment chooses a lockdown that is ex-post optimal but leads to ex-ante inefficient investment in expectation of the no-commitment outcome. This raises the possibility that constraints on government policy can prevent ex-ante inefficient policy outcomes.

5.1 Optimality of Limiting Future Policy Discretion

In our environment, a credible lockdown policy plan can be socially optimal. Suppose that rather than choosing a lockdown policy $L_t \in [0, 1]$ with discretion, the government is constrained to choosing a policy $L_t \in \mathcal{L}_t(\Omega_t) \subseteq [0, 1]$, where $\mathcal{L}_t(\Omega_t)$ is a subset of policies that depends on the prevailing health state Ω_t . As an example of a particularly heavy-handed policy constraint, consider $\mathcal{L}_t(\Omega_t) = \{L_t^c(\Omega_t)\}$. Then, the policy decision is constrained to the optimum under commitment, $L_t(\Omega_t) = L_t^c(\Omega_t)$. Clearly, this policy constraint implements the efficient outcome as it exactly mimics the time-consistent policy choice.

Going beyond this extreme example, we can study rules that constrain the extent of a lockdown. Consider a state-contingent rule $\mathcal{L}_t(\Omega_t) = \{L_t | L_t \leq \bar{L}_t(\Omega_t)\}$ so that a government at date t can choose any policy L_t that falls below $\bar{L}_t(\Omega_t)$ with discretion. In other words, the government commits to limiting the stringency of the lockdown.¹⁷ We then have the following result:

Proposition 2 (Value of Rules). *Consider a rule $\{\bar{L}_t(\Omega_t)\}_{t=0}^{\infty}$ such that $\bar{L}_t(\Omega_t) = L_t^c(\Omega_t)$ for all periods t and all health states Ω_t . Then, there exists an MPE subject to this rule, in which the government without commitment chooses the optimal policy under commitment.*

Proof. See Appendix A.2. □

Proposition 2 shows that the introduction of rules that impose a limit on the severity of lockdown can implement the optimal policy and therefore improve efficiency and welfare in an economy without government commitment. The idea behind the proof is as follows: Starting from the efficient policy sequence, a rule that takes the form of an upper bound allows only for downward

¹⁷This upper-bound rule is in line with Florida governor Ron DeSantis' announcement on September 25, 2020, of a statewide 50-percent-minimum capacity limit (i.e., an upper bound of 50 percent on capacity restrictions) for restaurants.

deviations in lockdown from L_t^c to some less strict lockdown $\tilde{L}_t < L_t^c$. But if a surprise relaxation of lockdown to the level \tilde{L}_t were optimal to a government without commitment given sunk investment $x^*(L_t^c, \Omega_t)$, which depends on the anticipated lockdown L_t^c and health state Ω_t , then a government with commitment could have implemented the same lockdown relaxation with firms anticipating it, leading to investment $x^*(\tilde{L}_t, \Omega_t)$. Since an anticipated lockdown relaxation yields higher investment and thus consumption, owing to the q-complementarity between investment and labor in production, such a deviation contradicts the optimality of the original lockdown policy under commitment.¹⁸

The intuition for this result is that an upper bound on lockdown stops the government without commitment from making short-sighted policy decisions when investment is treated as sunk. A lower bound on lockdown is not necessary, because lack of commitment is not associated with a temptation to impose too lax a lockdown. This is because q-complementarity between investment and labor in production (Assumption 1) implies that one-shot deviations from an equilibrium under commitment by a government without commitment are profitable only in the direction of stricter, not less strict, lockdown policy. For this reason, a lower bound on lockdown does not improve the efficiency of lockdown policy under the MPE considered in Proposition 2.

Note that on one hand, the rule described in Proposition 2 is less restrictive than one dictating the exact level of lockdown in every period and health state. On the other hand, an upper bound on lockdowns may still be overly strict if good reasons for imposing stricter lockdowns materialize in the future. While our analysis so far has abstracted from such reasons by assuming that the ex-post efficiency of future lockdowns can be guaranteed ex-ante, we now turn to a natural extension in which future policy flexibility is valuable.

5.2 Uncertainty and Noncontractible Information

Thus far, we have shown that under full information on the health state, a government without commitment would like to deviate from the optimal lockdown path and that rules limiting future lockdown can increase welfare by mitigating this commitment problem. In practice, of course, government policy depends not only on the health state but also on new information that arrives during a lockdown. Such information may include estimates of disease transmissibility and mor-

¹⁸We emphasize that our argument involves only the existence, not uniqueness, of an MPE that coincides with the efficient lockdown policy. In principle, there could exist other MPEs, but these would feature weakly lower welfare.

tality risk, the state of the economy, the likelihood of vaccine discovery, and the medical system's capacity. At the same time, information on future realizations of these variables may be hard to verify or to incorporate into a written contract.

This motivates us to study the design of rules under uncertainty and noncontractible information. We show that a modification of our previous result (Proposition 2) extends to an environment that incorporates such considerations. Specifically, we show that rules that constrain future government policy either as a function of future information revelation, as seen in the U.S. state of New York, or unconditionally, as in the U.S. state of Florida, can improve welfare.

To capture this idea, suppose that a state variable θ_t is realized, in addition to the prevailing health state Ω_t , before investment $x_t = x^*(L_t, \Omega_t, \theta_t)$ is made in anticipation of lockdown $L_t = L^*(x_t, \Omega_t, \theta_t)$ in period t . For simplicity, let θ_t be independently and identically distributed with associated probability density function $g(\theta_t)$ over support $[\underline{\theta}, \bar{\theta}]$ with $\underline{\theta} < \bar{\theta}$.¹⁹ Substituting the modified consumption function $c_t = c^*(x_t, L_t, \Omega_t, \theta_t)$ based on (8), social welfare at $t = 0$, given a sequence of state-contingent investment and lockdown policies $\{x_t(\Omega_t, \theta_t), L_t(\Omega_t, \theta_t)\}_{t=0}^{\infty}$, is

$$\begin{aligned} & \sum_{t=0}^{\infty} \beta^t \mathbb{E}_0[u(c^*(x_t, L_t, \Omega_t, \theta_t), \Omega_t, \theta_t)] & (21) \\ \text{s.t. } & \Omega_{t+1} = \Gamma(L_t, \Omega_t, \theta_t), \quad \forall t \geq 0, \\ & \theta_t \stackrel{iid}{\sim} g(\theta_t), \quad \forall t \geq 0, \\ & \Omega_0 \text{ given,} \end{aligned}$$

where the expectation $\mathbb{E}_0[\cdot]$ is taken over time-0 and future realizations of θ_t .

Note that the stochastic state θ_t enters the problem in multiple places. It indirectly enters the consumption function $c^*(\cdot)$ through its effect on production. At the same time, it directly enters the utility function $u(\cdot)$ and the SIRD model $\Gamma(\cdot)$. Finally, while equation (21) considers a given set of state-contingent investment and lockdown policies, the optimal investment function, $x^*(L_t, \Omega_t, \theta_t)$, and optimal lockdown function, $L^*(x_t, \Omega_t, \theta_t)$, also depend on θ_t .

In an MPE, the optimal lockdown policy under commitment depends on the health state Ω_t and the realization of θ_t , denoted by $L^c(\Omega_t, \theta_t)$. This policy function implicitly takes into account the optimal investment under commitment, $x^c(\Omega_t, \theta_t) = x^*(L^c(\Omega_t, \theta_t), \Omega_t, \theta_t)$. Analogously, the optimal lockdown policy in an MPE under lack of commitment depends only on the health state

¹⁹Our results are unchanged if the shock is persistent, assuming the shock is observable but not contractible.

Ω_t and the realization of θ_t , denoted by $L^n(\Omega_t, \theta_t)$. This policy function implicitly takes into account the MPE choice of investment under lack of commitment, $x^n(\Omega_t, \theta_t) = x^*(L_t^n(\Omega_t, \theta_t), \Omega_t, \theta_t)$.

Suppose that θ_t represents contractible information. Then, using an argument analogous to that in Proposition 2, it follows that a rule that imposes a sequence of upper bounds $\{\bar{L}_t(\Omega_t, \theta_t)\}_{t=0}^\infty$ on lockdown, so that $L_t \leq \bar{L}_t(\Omega_t, \theta_t)$, with $\bar{L}_t(\Omega_t, \theta_t) = L^c(\Omega_t, \theta_t)$ for all t , can increase social welfare by inducing the government without commitment to choose the policy under commitment.

In practice, some of the information in θ_t may not be contractible. In this case, a rigid plan may be too constraining, since policy flexibility in responding to realizations of θ_t is valuable. We show that bounded discretion in the form of a rule $\bar{L}_t(\Omega_t) > 0$ that constrains the government to policies $L_t \in [0, \bar{L}_t(\Omega_t)]$ independent of θ_t can still improve welfare in this case. To this end, consider the recursive formulation of the problem faced by a government without commitment:

$$W^n(x, \Omega, \theta) = \max_{L \in [0, 1]} \{u(c^*(x, L, \Omega, \theta), \Omega, \theta) + \beta \mathbb{E}_{\theta'}[V^n(\Gamma(L, \Omega, \theta), \theta')]\}, \quad (22)$$

$$\begin{aligned} V^n(\Omega', \theta') &= u(c^*(x^n(\Omega', \theta'), L^n(\Omega', \theta'), \Omega', \theta'), \Omega', \theta') \\ &\quad + \beta \mathbb{E}_{\theta''}[V^n(\Gamma(L^n(\Omega', \theta'), \Omega', \theta'), \theta'')], \end{aligned} \quad (23)$$

where $\Omega' = \Gamma(L, \Omega, \theta)$, and $\mathbb{E}_{\tilde{\theta}}[\cdot]$ denotes the expectation over the future realization of $\tilde{\theta}$. From here on, we operate under the following simplifying assumption:

Assumption 3. *The optimal lockdown policy under lack of commitment $L^n(\Omega_t, \theta_t)$ is strictly increasing in θ_t over interior $L^n(\Omega_t, \theta_t) \in (0, 1)$ and continuous in a neighborhood below $\bar{\theta}$ for all Ω_t . Moreover, the density $g(\cdot)$ is strictly positive and continuous in a neighborhood below $\bar{\theta}$.*

According to Assumption 3, higher values of the noncontractible state are associated with stricter optimal lockdown policies under lack of commitment. Then, we obtain the following result:

Proposition 3 (Value of Rules under Uncertainty). *Consider an MPE satisfying Assumption 3 for which lockdown policy is interior at time 0 for some realizations of θ_0 with positive probability. Then there exists a rule $\{\bar{L}_t(\Omega_t)\}_{t=0}^\infty$ and an MPE subject to this rule in which social welfare is strictly higher.*

Proof. See Appendix A.3. □

Proposition 3 shows that the introduction of rules increases social welfare, even if future policy discretion is valuable. The idea behind the proof is as follows: A government lacking commitment

chooses a more severe lockdown in the future than is socially desirable. Thus, a marginally binding cap on lockdowns increases social welfare by raising investment and therefore output and consumption at no efficiency cost. To arrive at this conclusion, a key part of the argument is that the most extreme lockdown policy imposed by the government without commitment is never optimal for a government with commitment under any realization of new information. This is natural in our setting in which the production technology satisfies an Inada condition—completely shutting down the economy yields unbounded marginal gains from opening the economy slightly.²⁰

The intuition for this result is that a marginally binding rule does not prevent efficient lockdowns while limiting the damages of excessive lockdowns in the future. By preventing only the most extreme variants of future lockdown policies, such a rule can improve the efficiency of firms' investment choice and thereby increase social welfare.

6 Quantitative Exercise

We now illustrate the quantitative implications of lack of government commitment during a pandemic in an illustrative calibration and simulation of our model. The goal is to compare the lockdown policy, aggregate output and consumption, and the health state in an economy with a pandemic subject to the efficient lockdown policy under commitment versus the inefficient lockdown policy under no commitment. This comparison also allows us to illustrate how rules that limit lockdown discretion, which we have shown to be associated with efficient lockdown policy (Proposition 2), affect the path of a pandemic.

6.1 Calibration

In order to calibrate our model, we make several assumptions about the production technology, the SIRD model of disease spread, and preferences. The main steps of our calibration strategy are outlined here; further details are in Appendix B.

We start by specifying the production technology. We assume that gross output, y_t , is generated according to a Cobb-Douglas production function that combines investment, x_t , with labor,

²⁰As an example of a case in which extreme choices are sometimes optimal even under commitment, see [Halac and Yared \(2020\)](#) for a discussion of threshold contracts with escape clauses.

ℓ_t , given by

$$y_t = Ax_t^\alpha \ell_t^{1-\alpha}, \quad (24)$$

where A is total factor productivity and $\ell_t = (1 - L_t)(S_t + \gamma I_t + R_t)$ is the effective labor input. Effective labor input may be less than the unit mass of the initial population owing either to deaths from past infections or to the lower relative productivity of infected workers, indexed by $\gamma \in [0, 1]$. The assumption of a Cobb-Douglas production function implies that the total wage bill $w_t \ell_t$, and hence aggregate consumption c_t , is a constant share $(1 - \alpha)$ of gross output y_t . We set $\gamma = 0.5$; that is, we assume that the infected subpopulation works at 50 percent capacity, roughly corresponding to the share of asymptomatic infections according to [Yanes-Lane et al. \(2020\)](#).²¹

Our choice of α is based on the factor share of intermediate inputs x_t that are subject to the type of time inconsistency problem described in our model. To determine the appropriate value of α , we focus on intermediate inputs in U.S. input-output tables satisfying all three of the following criteria. First, we require inputs to be typically purchased in advance and therefore chosen in anticipation of future lockdown policy. Second, we require inputs to be such that reimbursement in the event of a surprise lockdown is unlikely. Finally, we require inputs to be perishable or not easily storable so that mistakenly purchasing them with wrong expectations of future lockdown is costly. Using the intersection of these three criteria, we find that the cost of intermediate inputs corresponding to investment x_t in our model makes up 51.6 percent of the cost of all intermediate inputs and 78.2 percent of the cost of compensation of employees.²² We set $\alpha = 0.439$ to match the ratio of the cost of intermediate inputs corresponding to investment x_t in our model relative to the sum of the cost of these intermediate inputs and the cost of compensation of employees, which equals 43.9 percent. We acknowledge that this value of α may be imprecise, since not all inputs within an input category can be perfectly included or excluded from our criteria. For this reason, we consider the robustness of our quantitative results to the value of α .

Next, we specify the SIRD model of disease spread. To this end, we set the period length equal

²¹While we are not aware of any direct evidence on individual workers' productivity throughout the disease stages, we introduce this parameter, γ , in order to allow for the possibility that the pandemic may have a direct effect on the efficiency of the economy. Extensive simulations indicate that our results are not particularly sensitive to this parameter.

²²See Tables A–E in the Online Appendix for details of the classification of intermediate inputs by six-digit BEA industry code.

to one week. The health state $\Omega_t = \{S_t, I_t, R_t, D_t\}$ obeys the following law of motion:

$$S_{t+1} = [1 - (\rho_{1,t}(1 - L_t)^2 + \rho_{2,t}) I_t] S_t \quad (25)$$

$$I_{t+1} = (1 - \rho_3 - \rho_4) I_t + (\rho_{1,t}(1 - L_t)^2 + \rho_{2,t}) I_t S_t \quad (26)$$

$$R_{t+1} = R_t + \rho_3 I_t \quad (27)$$

$$D_{t+1} = D_t + \rho_4 I_t. \quad (28)$$

The intuition behind equations (25)–(28) is as follows. The total mass of new infections corresponds to a flow from the current susceptible state, S_t , to next period's infected state, I_{t+1} . New infections obtain as a result of infected individuals meeting susceptible individuals, either at or outside of work. Specifically, a fraction $\rho_{1,t}$ of all meetings between $(1 - L_t)I_t$ infected workers and $(1 - L_t)S_t$ susceptible workers result in disease transmission at work, while a fraction $\rho_{2,t}$ of all meetings between I_t infected individuals and S_t susceptible individuals result in disease transmission outside of work. Therefore, the total flow from the current susceptible state, S_t , to next period's infected state, I_{t+1} , is given by $(\rho_{1,t}(1 - L_t)^2 + \rho_{2,t})I_t S_t$. At the same time, a fraction ρ_3 of currently infected individuals I_t recover and become part of the state variable R_{t+1} next period, while a fraction ρ_4 of currently infected individuals I_t pass away and become part of the state variable D_{t+1} next period. Based on the SIRD model in equations (25)–(28), the basic reproduction number is $\mathcal{R}_0 = (\rho_{1,0} + \rho_{2,0})/(\rho_3 + \rho_4)$, which corresponds to the number of new infections per infected individual in the early stage of the pandemic.

That $\rho_{1,t}$ and $\rho_{2,t}$ are allowed to depend on time t reflects the fact that the arrival of a vaccine may affect these transition rates. Specifically, denoting by $T > 0$ the deterministic date of arrival of a vaccine, which is assumed to eliminate any further disease transmission, we let

$$\rho_{1,t} = \begin{cases} \rho_1 & \text{for } t < T, \\ 0 & \text{for } t \geq T, \end{cases}, \quad \rho_{2,t} = \begin{cases} \rho_2 & \text{for } t < T, \\ 0 & \text{for } t \geq T, \end{cases} \quad (29)$$

for some fixed values ρ_1 and ρ_2 .

The SIRD model in equations (25)–(28) is fully parameterized by the vector $[\rho_{1,t}, \rho_{2,t}, \rho_3, \rho_4]$, which we discipline using empirical evidence on disease transitions associated with COVID-19. Specifically, we calibrate our model using the following set of equations that relate functions of

model parameters to empirical moments of the data:

$$\text{average length of infection in weeks: } \frac{1}{\rho_3 + \rho_4} = 2.000; \quad (30)$$

$$\text{mortality rate conditional on infection: } \frac{\rho_4}{\rho_3 + \rho_4} = 0.058; \quad (31)$$

$$\text{basic reproduction number } \mathcal{R}_0: \frac{\rho_1 + \rho_2}{\rho_3 + \rho_4} = 1.660. \quad (32)$$

We choose these target moments based on recent scientific evidence on disease dynamics of the SARS-CoV-2 virus. Specifically, we adopt an average length of infection of two weeks, following recent guidelines by health officials ([Centers for Disease Control and Prevention, 2021](#)). The mortality rate conditional on infection is taken as the peak mortality rate following March 7, 2020, which is after the large initial spike ([COVID Tracking Project, 2021](#)). The basic reproduction number \mathcal{R}_0 is the median among United States counties according to [Sy et al. \(2021\)](#).

In addition to the three equations (30)–(32), we assume that the probability of infection when working is 50.0 percent higher than when not working:

$$\rho_1 + \rho_2 = 1.5\rho_2. \quad (33)$$

All of these parameters are calibrated for an economy without any lockdown, so the values of the average length of infection, the conditional mortality rate, and basic reproduction number \mathcal{R}_0 correspond roughly to the early stage of the pandemic in the first quarter of 2020 in the United States. Together, equations (30)–(33) yield the following set of calibrated SIRD model parameters:

$$\rho_1 = 0.277 \quad (34)$$

$$\rho_2 = 0.553 \quad (35)$$

$$\rho_3 = 0.471 \quad (36)$$

$$\rho_4 = 0.029. \quad (37)$$

Finally, we turn to specifying preferences. We assume that period utility is additively separable between log utility over per-capita consumption, $\bar{c}_t = c_t / (S_t + I_t + R_t)$, and a flow value of being

alive, v , with the value of being dead normalized to zero:

$$u(c_t, \Omega_t) = (S_t + I_t + R_t) [\ln(\bar{c}_t) + v]. \quad (38)$$

Lifetime utility is simply the discounted stream of period utilities $\{u(\bar{c}_t, \Omega_t)\}_{t \geq 0}$ with period discount factor β . We set the flow value of being alive, $v = 4.545$, which corresponds to a value of a statistical life of USD 11.5 million (Greenstone and Nigam, 2020; Glover et al., 2020). We choose the weekly interest rate r to match an annual interest rate of 3.0 percent and the weekly discount factor β such that $\beta(1 + r) = 1$:

$$r = 1.03^{1/52} - 1 \approx 5.686 \times 10^{-4} \quad (39)$$

$$\beta = \frac{1}{1 + r} \approx 9.994 \times 10^{-1}. \quad (40)$$

In all simulations, we assume that the economy starts out with a population of 331 million agents, out of which all are susceptible, except for 100 initial infections. In our baseline calibration, we assume that at time $T = 52$ a vaccine arrives, which ends the possibility of new infections occurring for all $t \geq T$.

Table 1 summarizes our calibration of the model's parameters.

6.2 Model Simulations

In order to simulate the economy with and without commitment, we use backward induction to solve the problem of the government. We first compute the continuation value of reaching period T , in which a vaccine becomes available. From this period onwards, there is no commitment problem, since forgoing lockdown is always optimal. We then solve the model backward for $t = T - 1, T - 2, \dots, 0$.

Figure 1 compares the optimal policy under commitment to that under lack of commitment. The results are consistent with our theoretical predictions: Lockdown under lack of commitment is more severe than lockdown under commitment. Panel (a) shows that compared with the economy under commitment, lockdown is more severe at most points in time under lack of commitment. Panel (b) illustrates the consequences for aggregate consumption, which is proportional to gross output and declines significantly more under lack of commitment. Panels (c)–(f) display the con-


Table 1. Model Calibration Results

Parameter	Description	Value	Target
<i>Panel A. Technology parameters</i>			
α	Cobb-Douglas parameter	0.439	Intermediate-input share of 43.9 percent
γ	Rel. prod. of infected workers	0.500	Asympt. infections share of 50.0 percent
r	Weekly interest rate	0.001	Annual interest rate of 3.0 percent
<i>Panel B. SIRD model of disease spread parameters</i>			
ρ_1	At-work infection rate	0.277	Basic reproduction number \mathcal{R}_0 of 1.660
ρ_2	Not-at-work infection rate	0.553	50.0 percent higher infection risk at work
ρ_3	Recovery rate	0.471	Average length of infection of 14 days
ρ_4	Death rate	0.029	Mortality rate of 5.8 percent
S_0	Initial susceptible share	> 0.999	$1 - 100/331,000,000$ initially susceptible
I_0	Initial infected share	< 0.001	$100/331,000,000$ initially infected
R_0	Initial recovered share	0.000	No initially recovered
D_0	Initial dead share	0.000	No initially dead
T	Vaccine arrival time	52	Arrives 1 year after start of pandemic
<i>Panel C. Preference parameters</i>			
ν	Value of life	4.545	Value of statistical life of USD 11.5mm
β	Discount factor	0.999	$\beta(1+r) = 1$

Notes: This table shows the calibrated model parameters along with the corresponding empirical target moments. See text for details.

sequences of lack of commitment for health outcomes. Because lockdown is more severe under lack of commitment, fewer individuals are exposed to the disease. Consequently, the share of the population that is susceptible at any point in time is higher, the share infected is lower, the share recovered is lower, and the share dead is lower.


Figure 1. Simulations with versus without Commitment



Notes: This figure shows time series of lockdown share L_t (panel a), aggregate consumption c_t (panel b), the share of susceptibles S_t (panel c), the share of infected I_t (panel d), the share of recovered R_t (panel e), and the share of dead D_t (panel f). All are plotted against weeks since the outbreak of the pandemic. The red short-dashed line with diamonds represents outcomes under lockdown policy with commitment, while the green dash-dotted line with triangles represents outcomes under lockdown policy without commitment. The vertical striped line at week 52 indicates the arrival of a vaccine. See text for details.

We examine how these results are affected by the value of the intermediate input share (α), which we set to 0.439 in our benchmark calibration. Figure 2 panel (a) shows the paths of aggregate consumption under a low factor share value of $\alpha = 0.220$, while panel (b) shows the paths under the benchmark calibration with $\alpha = 0.439$, and panel (c) shows the path under a high factor share value of $\alpha = 0.878$. According to this figure, the results under either low or high factor share values α are qualitatively consistent with the results from the benchmark model. In all three cases, lockdown under lack of commitment is more severe, leading to larger consumption losses than those under commitment. Quantitatively, both the severity of lockdown under either policy regime and the discrepancy between the policy with versus without commitment are increasing in the intermediate-input factor share α . Nevertheless, even under a low factor share value of α corresponding to half that in our benchmark calibration (panel a), we find a significant discrepancy between lockdown with versus without commitment.

Figure 2. Time Series of Aggregate Consumption under Different Factor Shares



Notes: This figure shows time series of aggregate consumption c_t under different factor shares: a low factor share ($\alpha = 0.220$, panel a), the benchmark factor share ($\alpha = 0.439$, panel b), and a high factor share ($\alpha = 0.878$, panel c). The red short-dashed line with diamonds represents outcomes under lockdown policy with commitment, while the green dash-dotted line with triangles represents outcomes under lockdown policy without commitment. The vertical striped line at week 52 indicates the arrival of a vaccine. See text for details.

6.3 Comparative Statics

We now examine how excessively severe lockdown due to lack of commitment depends on features of the economic environment. Table 2 considers the consumption loss during the first year of the pandemic due to lack of commitment for different parameter values. The first two columns show the consumption loss under commitment and under lack of commitment relative to an economy without a pandemic, while the third column shows the consumption loss under lack of commitment relative to commitment. As shown in the first row, our calibrated benchmark economy

predicts that lack of commitment reduces aggregate consumption by 14.9 percent. These findings suggest that optimal policy commitments to limit lockdown could result in a significant reduction of consumption losses during a pandemic.

The subsequent rows show comparative statics with respect to a low value (i.e., half the value in our benchmark) and a high value (i.e., double the value in our benchmark) for each of six model parameters. When we consider different values for the intermediate input share (α), a higher intermediate input share is associated with greater consumption losses due to lack of commitment. These results are consistent with the larger gaps in consumption between the commitment case and lack of commitment case in Figure 2 panel (b) relative to Figure 2 panel (a). Intuitively, investment distortions due to lack of commitment are more impactful for higher values of the intermediate input share. When we consider different values of the discount rate ($1 - \beta$), we find that the consumption loss due to lack of commitment is larger for lower discount rates. This is intuitive: The more the government values the future, the larger the temptation to renege on past promises to limit lockdown, since the perceived benefits of mitigating future disease spread are larger. A similar intuition explains why the consumption loss due to lack of commitment is larger if the value of life (ν) is larger, since the government without commitment overweighs the value of life relative to the efficient solution. Moreover, the higher the transmission rate of disease at work (ρ_1), the more beneficial is lockdown on the margin, and the larger the temptation to renege on past promises for a limited lockdown, and thus the larger the consumption loss due to lack of commitment. An analogous reasoning explains why the consumption loss due to lack of commitment is larger if the transmission rate outside of work (ρ_2) is higher, since in that case, mitigating transmission at work through lockdown can further reduce transmission outside of work. Finally, the consumption loss due to lack of commitment is larger for longer times until vaccine arrival (T), because a longer waiting period increases the duration of the commitment problem.

7 Concluding Remarks

We have analyzed the value of government commitment in designing lockdown policies. In our model, a government would like to commit to limit the extent of future lockdowns in order to support more optimistic expectations and stimulate investment in the present. However, such a commitment is not credible, since investment decisions are sunk when the government makes the

Table 2. Aggregate Consumption Loss during First Year of Pandemic

		Aggregate consumption loss (%)		
		C	NC	NC vs. C
Baseline		32.0	42.6	15.7
Intermediate input share, α	Low	29.8	38.3	12.1
	High	42.0	59.1	29.6
Discount rate, $1 - \beta$	Low	35.9	46.0	15.8
	High	28.2	39.3	15.4
Value of life, ν	Low	30.8	41.5	15.5
	High	33.9	44.2	15.7
Transmission rate at work, ρ_1	Low	4.7	18.6	14.6
	High	53.5	61.4	17.0
Transmission rate outside of work, ρ_2	Low	0.0	0.6	0.6
	High	23.4	28.1	6.1
Vaccine arrival time, T	Low	0.2	4.4	4.2
	High	52.5	55.8	6.9

Notes: This table shows aggregate consumption losses in percentage points, calculated by summing over aggregate consumption during the first 52 weeks of the pandemic, discounted at a weekly interest rate that corresponds to an annual compound interest rate of 3.0 percent. The rightmost three columns report and compare two economies: One with lockdown policy under commitment (C) and one with lockdown policy under no commitment (NC), both relative to the economy without a pandemic. The third column (NC vs. C) contains the aggregate consumption loss from no commitment relative to that under commitment. The “baseline” results are those obtained using the calibrated model. For the two economies and their comparison, comparative statics in each of six model parameters are conducted: The intermediate input share (α), the discount rate ($1 - \beta$), the value of life (ν), the transition rate of infections at work (ρ_1), the transition rate of infections outside of work (ρ_2), and the vaccine arrival time (T). For each parameter of the comparative statics, results are shown for a “low” value of half the calibrated baseline parameter and a “high” value of twice the calibrated baseline parameter. See text for details.

lockdown decision. This gives value to rules limiting future lockdown policy discretion. We illustrate the distortions introduced by lack of commitment and its comparative statics with respect to fundamental model parameters in a quantitative exercise using a calibrated version of our model.

Our analysis points to several interesting avenues for future research. First, the generality of our approach suggests that time consistency considerations could be relevant to many lockdown decision problems. For instance, it would be interesting to characterize the optimal policy response to widespread employee furloughs. Payroll subsidies and cheap access to credit for businesses have been widely advocated during the global COVID-19 pandemic. However, their efficiency under lack of government commitment could be drastically different from that under commitment, which previous work has focused exclusively on. Time inconsistency is also relevant in other domains such as school and college decisions to reopen in anticipation of future

lockdowns or private investments in disease-mitigating equipment. Insights similar to our characterization of lockdown policy under lack of commitment may apply in such contexts.

Second, our evaluation of the effect of rules that limit lockdowns assumes that governments adhere to such rules. In practice, rules may be broken, and the private sector may be uncertain about the government's commitment to respecting them. In the context of capital taxation, [Phelan \(2006\)](#) and [Dovis and Kirpalani \(2019\)](#) show that this consideration leads the private sector to dynamically update its beliefs about a government's ability to commit. We conjecture that in our framework, this uncertainty could cause firms to react to lockdown extensions by becoming increasingly pessimistic about the government's ability to commit to lifting a future lockdown. This could lead to further declines in investment and economic activity, as well as political economy consequences of lockdown extensions.

Finally, our analysis ignores the availability of monetary and fiscal policy tools, which are considered in contemporaneous work by [Guerrieri et al. \(2020\)](#). In our framework, these tools could not only mitigate the immediate economic costs of a pandemic but also boost investment, thus counteracting future economic costs from underinvestment due to the government's lack of commitment. We leave the exploration of how optimal lockdown policy interacts with monetary and fiscal policy under lack of government commitment as an interesting subject of further research.

References

- Acemoglu, Daron, Victor Chernozhukov, Iván Werning, and Michael D. Whinston, "Optimal Targeted Lockdowns in a Multi-Group SIR Model," *Working Paper*, 2020.
- Aguiar, Mark and Manuel Amador, "Growth in the Shadow of Expropriation," *Quarterly Journal of Economics*, 2011, 126 (2), 651–697.
- , –, and Gita Gopinath, "Investment Cycles and Sovereign Debt Overhang," *Review of Economic Studies*, 2009, 76 (1), 1–31.
- Alvarez, Fernando E., David Argente, and Francesco Lippi, "A Simple Planning Problem for COVID-19 Lockdown," *NBER Working Paper No. 26981*, 2020.
- Amador, Manuel, Iván Werning, and George-Marios Angeletos, "Commitment vs. Flexibility," *Econometrica*, 2006, 74 (2), 365–396.
- Athey, Susan, Andrew Atkeson, and Patrick J. Kehoe, "The Optimal Degree of Discretion in Monetary Policy," *Econometrica*, 2005, 73 (5), 1431–1475.
- Atkeson, Andrew, "Lockdowns and GDP: Is There a Tradeoff?," *Working Paper*, 2020.
- , "What Will Be the Economic Impact of COVID-19 in the US? Rough Estimates of Disease Scenarios," *NBER Working Paper No. 26867*, 2020.
- Baker, Scott R., Nicholas Bloom, Steven J. Davis, and Stephen J. Terry, "COVID-Induced Economic Uncertainty," *NBER Working Paper No. 26983*, 2020.
- Berger, David W., Kyle F. Herkenhoff, and Simon Mongey, "An SEIR Infectious Disease Model with Testing and Conditional Quarantine," *NBER Working Paper No. 26901*, 2020.
- Centers for Disease Control and Prevention, "Interim Guidance on Ending Isolation and Precautions for Adults with COVID-19," <https://www.cdc.gov/coronavirus/2019-ncov/hcp/duration-isolation.html> (accessed on July 23, 2021), 2021.
- Chari, V. V. and Patrick J. Kehoe, "Sustainable Plans," *Journal of Political Economy*, 1990, 98 (4), 783–802.
- , Juan Pablo Nicolini, and Pedro Teles, "Optimal Capital Taxation Revisited," *Journal of Monetary Economics*, 2019.
- , Rishabh Kirpalani, and Christopher Phelan, "The Hammer and the Scalpel: On the Economics of Indiscriminate versus Targeted Isolation Policies during Pandemics," *NBER Working Paper No. 27232*, 2020.
- , –, and Luis Perez, "On the Efficiency of Competitive Equilibria with Pandemics," *Working Paper*, 2021.
- COVID Tracking Project, "The Data: Totals for the US," <https://covidtracking.com/data/national> (accessed on July 23, 2021), 2021.
- DeSantis, Ron, "Ron DeSantis on the Pandemic Year: Don't Trust the Elites," *Wall Street Journal* (accessed on July 23, 2021), 2021.
- Dingel, Jonathan I. and Brent Neiman, "How Many Jobs Can Be Done at Home?," *NBER Working Paper No. 26948*, 2020.
- Dovis, Alessandro and Rishabh Kirpalani, "Rules without Commitment: Reputation and Incentives," *NBER Working Paper No. 26451*, 2019.
- Eichenbaum, Martin S., Sergio Rebelo, and Mathias Trabandt, "The Macroeconomics of Epidemics," *NBER Working Paper No. 26882*, 2020.
- , –, and –, "The Macroeconomics of Testing and Quarantining," *Working Paper*, 2020.
- Farboodi, Maryam, Gregor Jarosch, and Robert Shimer, "Internal and External Effects of Social Distancing in a Pandemic," *Working Paper*, 2020.
- Ferguson, Neil M., Daniel Laydon, Gemma Nedjati-Gilani, Natsuko Imai, Kylie Ainslie, Marc Baguelin, Sangeeta Bhatia, Adhiratha Boonyasiri, Zulma Cucunubá, Gina Cuomo-Dannenburg, Amy Dighe, Iliaria Dorigatti, Han Fu, Katy Gaythorpe, Will Green, Arran Hamlet, Wes Hinsley, Lucy C. Okell, Sabine van Elsland, Hayley Thompson, Robert Verity, Erik Volz, Haowei Wang, Yuanrong Wang, Patrick G. T. Walker, Caroline Walters, Peter Winskill, Charles Whittaker, Christl A. Donnelly, Steven Riley, and Azra C. Ghani, "Impact of Non-

- Pharmaceutical Interventions (NPIs) to Reduce COVID-19 Mortality and Healthcare Demand," *Working Paper*, 2020.
- Glover, Andrew, Jonathan Heathcote, Dirk Krueger, and José-Víctor Ríos-Rull**, "Health versus Wealth: On the Distributional Effects of Controlling a Pandemic," *Working Paper*, 2020.
- Gourinchas, Pierre-Olivier**, "Flattening the Pandemic and Recession Curves," *Working Paper*, 2020.
- Greenstone, Michael and Vishan Nigam**, "Does Social Distancing Matter?," *University of Chicago, Becker Friedman Institute for Economics Working Paper No. 2020-26*, 2020.
- Guerrieri, Veronica, Guido Lorenzoni, Ludwig Straub, and Iván Werning**, "Macroeconomic Implications of COVID-19: Can Negative Supply Shocks Cause Demand Shortages?," *NBER Working Paper No. 26918*, 2020.
- Halac, Marina and Pierre Yared**, "Fiscal Rules and Discretion under Persistent Shocks," *Econometrica*, 2014, 82 (5), 1557–1614.
- and – , "Fiscal Rules and Discretion in a World Economy," *American Economic Review*, 2018, 108 (8), 2305–2334.
- and – , "Commitment versus Flexibility with Costly Verification," *Journal of Political Economy*, 2020, 128 (12), 4523–4573.
- Hall, Robert E., Charles I. Jones, and Peter J. Klenow**, "Trading Off Consumption and COVID-19 Deaths," *Working Paper*, 2020.
- Hassan, Tarek A., Stephan Hollander, Laurence van Lent, Markus Schwedeler, and Ahmed Tahoun**, "Firm-Level Exposure to Epidemic Diseases: Covid-19, SARS, and H1N1," *NBER Working Paper No. 26971*, 2020.
- Kermack, William O. and Anderson G. McKendrick**, "A Contribution to the Mathematical Theory of Epidemics," *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 1927, 115 (772), 700–721.
- Klein, Paul, Per Krusell, and José-Víctor Ríos-Rull**, "Time-Consistent Public Policy," *Review of Economic Studies*, 2008, 75 (3), 789–808.
- Kydland, Finn E. and Edward C. Prescott**, "Dynamic Optimal Taxation, Rational Expectations and Optimal Control," *Journal of Economic Dynamics and Control*, 1980, 2, 79–91.
- Mongey, Simon, Laura Pilossoph, and Alex Weinberg**, "Which Workers Bear the Burden of Social Distancing Policy?," *Working Paper*, 2020.
- Moser, Christian and Pedro Olea de Souza e Silva**, "Optimal Paternalistic Savings Policies," *Working Paper*, 2019.
- Phelan, Christopher**, "Public Trust and Government Betrayal," *Journal of Economic Theory*, 2006, 130 (1), 27–43.
- Smit, Sven, Martin Hirt, Penny Dash, Audrey Lucas, Tom Latkovic, Matt Wilson, Ezra Greenberg, Kevin Buehler, and Klemens Hjartar**, "Crushing Coronavirus Uncertainty: The Big 'Unlock' for Our Economies," *Report, McKinsey & Company*, 2020.
- Sy, Karla Therese L., Laura F. White, and Brooke E. Nichols**, "Population Density and Basic Reproductive Number of COVID-19 across United States Counties," *PLOS ONE*, 2021, 16 (4), 1–11.
- U.S. Bureau of Economic Analysis**, "Input-Output Accounts Data," <https://www.bea.gov/industry/input-output-accounts-data> (accessed on July 23, 2021), 2021.
- Yanes-Lane, Mercedes, Nicholas Winters, Federica Fregonese, Mayara Bastos, Sara Perlman-Arrow, Jonathon R. Campbell, and Dick Menzies**, "Proportion of Asymptomatic Infection among COVID-19 Positive Persons and their Transmission Potential: A Systematic Review and Meta-Analysis," *PLOS ONE*, 2020, 15 (11), 1–21.

Appendix

A Proofs

A.1 Proof of Proposition 1

Proof. To prove that the optimal lockdown policy is time inconsistent, we want to show that $L_t^c \neq L_t^n$ for some t . Let t be a period in which $L_t^c \in (0, 1)$, which exists by assumption. Suppose, by way of contradiction, that there exists an MPE under no commitment that coincides with the optimal policy under commitment in all possible states and all periods. For a government choosing lockdown L_t given health state Ω_t , this would mean that the continuation value would be the same with and without commitment. Therefore,

$$\frac{dV^c(\cdot)}{dL_t} = \frac{dV^n(\cdot)}{dL_t}, \quad (41)$$

meaning that the derivative of the continuation value with respect to current lockdown is the same with and without commitment. However, if (41) holds, then the optimality condition of the government with commitment in (13) and that of the government without commitment in (18) cannot simultaneously hold because

$$-\beta \frac{dV^c(\cdot)}{dL_t} = \frac{\partial u(\cdot)}{\partial c_t} \left[\frac{\partial c^*(\cdot)}{\partial x_t} \frac{\partial x^c(\cdot)}{\partial L_t} + \frac{\partial c^*(\cdot)}{\partial L_t} \right] < \frac{\partial u(\cdot)}{\partial c_t} \frac{\partial c^*(\cdot)}{\partial L_t} = -\beta \frac{dV^n(\cdot)}{dL_t}, \quad (42)$$

where the strict inequality follows from Assumption 1. This poses a contradiction with the equality in (41), proving the claim that the policy under lack of commitment does not coincide with that under commitment. Therefore, the optimal lockdown policy is time inconsistent. \square

A.2 Proof of Proposition 2

Proof. To prove that a rule consisting of an upper bound $\bar{L}_t(\Omega_t) = L_t^c(\Omega_t)$ on L_t supports an MPE that attains the efficient allocation, we want to show that there exists no profitable deviation from this allocation by a government without commitment adhering to this rule. Consider a government today choosing lockdown policy under the efficient state-contingent rule and expecting all future governments to choose lockdown equal to the efficient state-contingent rule. Therefore, the

government's state-contingent policy is given by $\{L_t(\Omega_t)\}_{t=0}^{\infty}$ such that $L_t(\Omega_t) = \bar{L}_t(\Omega_t) = L_t^c(\Omega_t)$ in all states and all periods, which induces a sequence of investments $\{x_t^c(\Omega_t)\}_{t=0}^{\infty}$ such that $x_t(\Omega_t) = x_t^c(\Omega_t)$ in all states and all periods. Now consider in any period t the problem of the government without commitment, which anticipates that all future governments will follow the optimal policy under commitment and also investment will match that under commitment. Given all this, when we compare the FOC of the government under lack of commitment (18) with that under commitment (13), the unconstrained government without commitment would like to choose a value of L_t that is strictly higher than $\bar{L}_t(\Omega_t)$. Clearly, this is not possible given the rule, which constrains the government to choose $L_t \leq \bar{L}_t(\Omega_t)$. Thus, there are two possibilities: Either $\bar{L}_t(\Omega_t) = L_t^c(\Omega_t) > 0$ and there exists a profitable downward deviation to some $\tilde{L}_t \in [0, L_t^c(\Omega_t))$ in period t , or else the current allocation constitutes an MPE. Suppose by way of contradiction there exists such a profitable downward deviation from $L_t^c(\Omega_t) > 0$ to $\tilde{L}_t < L_t^c(\Omega_t)$ in period t given sunk investment $x_t(\Omega_t)$ and health state Ω_t . For this to be the case, we must have

$$\begin{aligned} & u(c^*(x^*(L_t^c(\Omega_t), \Omega_t), \tilde{L}_t, \Omega_t)) + \beta V^c(\Gamma(\tilde{L}_t, \Omega_t)) \\ & > u(c^*(x^*(L_t^c(\Omega_t), \Omega_t), L_t^c(\Omega_t), \Omega_t)) + \beta V^c(\Gamma(L_t^c(\Omega_t), \Omega_t)). \end{aligned} \quad (43)$$

Because this deviation is unanticipated, investment $x^*(L_t^c(\Omega_t), \Omega_t)$ remains at the level in expectation of lockdown $L_t^c(\Omega_t)$ under any deviation of investment \tilde{L}_t . We now show that if the inequality in (43) were to hold, then the government under commitment could profitably deviate from its investment strategy, thus contradicting the optimality of the original MPE. Consider the same deviation from $L_t^c(\Omega_t) > 0$ to $\tilde{L}_t < L_t^c(\Omega_t)$ by a government with commitment. Since firms would anticipate this new lockdown policy in period t under commitment, q-complementarity between x_t and ℓ_t in production (Assumption 1) implies that the optimal investment would also adjust upward from $x_t = x^*(L_t^c(\Omega_t), \Omega_t)$ to $\tilde{x}_t = x^*(\tilde{L}_t, \Omega_t) > x_t$. Since consumption in (8) is strictly increasing in x_t , this deviation yields a strictly greater benefit to the government with commitment compared with that of the government under commitment. We conclude that equation (43) characterizing the deviation by the government without commitment can hold only if there exists a profitable deviation by the government with commitment. This contradicts the optimality of the original MPE, thus invalidating the existence of a profitable downward deviation by the government without commitment. Therefore, the allocation under commitment, together with a

rule consisting of an upper bound $\bar{L}_t(\Omega_t) = L_t^c(\Omega_t)$ on $L_t(\Omega)$ in all states and all periods, also constitutes an MPE under lack of commitment. \square

A.3 Proof of Proposition 3

Proof. First, note that lockdown under full commitment and under lack of commitment are never maximal, owing to the Inada condition on the production function $f(\cdot)$ with respect to labor input ℓ .

Since the statement of the proposition concerns the existence of a rule in some period t , we will consider period $t = 0$. Now contemplate a rule that imposes an upper bound $\bar{L}(\Omega_0; \varepsilon) = L^n(\Omega_0, \bar{\theta} - \varepsilon)$, for some $\varepsilon > 0$, on labor supply L_0 at time 0 given Ω_0 . We will establish that such a rule strictly increases social welfare for small enough $\varepsilon > 0$. For the remainder of the proof, we consider a perturbation only at time $t = 0$, which we treat as the current period, and will drop all time subscripts.

For a given state (Ω, θ) , let $x^n \equiv x^n(\Omega, \theta)$ and $L^n(\Omega, \theta)$ denote the MPE investment policy and lockdown policy under no commitment in the absence of a rule, and let $x^r \equiv x^r(\Omega, \theta; \varepsilon)$ and $L^r(\Omega, \theta; \varepsilon)$ denote the MPE investment policy and lockdown policy under no commitment subject to the rule $\bar{L}(\Omega; \varepsilon)$, all from a period-0 perspective. Now let us look at the welfare in an economy subject to such a rule relative to that in an economy without rules. By Assumption 3, $L^n(\Omega, \theta)$ is strictly increasing in θ , so the difference in social welfare between lockdown with or without the rule is zero conditional on $\theta < \bar{\theta} - \varepsilon$, since the policy under no commitment is unaffected by the rule for these realizations of θ . The difference in social welfare from realizations $\theta \in [\bar{\theta} - \varepsilon, \bar{\theta}]$ is nonzero and equals

$$\int_{\theta=\bar{\theta}-\varepsilon}^{\bar{\theta}} \left\{ \begin{aligned} & [u(c^*(x^r, L^r(\Omega, \theta; \varepsilon), \Omega, \theta), \Omega, \theta) + \beta \mathbb{E}_{\theta'}[V^n(\Gamma(L^r(\Omega, \theta; \varepsilon), \Omega, \theta), \theta')]] \\ & - [u(c^*(x^n, L^n(\Omega, \theta), \Omega, \theta), \Omega, \theta) + \beta \mathbb{E}_{\theta'}[V^n(\Gamma(L^n(\Omega, \theta), \Omega, \theta), \theta')]] \end{aligned} \right\} g(\theta) d\theta, \quad (44)$$

where $\mathbb{E}_{\theta'}[\cdot]$ denotes the current period's expectation over next period's realization of θ' . We first

establish that (44) is bounded from below by

$$\int_{\theta=\bar{\theta}-\varepsilon}^{\bar{\theta}} \left\{ \begin{aligned} & [u(c^*(x^n, L^n(\Omega, \bar{\theta} - \varepsilon), \Omega, \theta), \Omega, \theta) + \beta \mathbb{E}_{\theta'}[V^n(\Gamma(L^n(\Omega, \bar{\theta} - \varepsilon), \Omega, \theta), \theta')]] \\ & - [u(c^*(x^n, L^n(\Omega, \theta), \Omega, \theta), \Omega, \theta) + \beta \mathbb{E}_{\theta'}[V^n(\Gamma(L^n(\Omega, \theta), \Omega, \theta), \theta')]] \end{aligned} \right\} g(\theta) d\theta, \quad (45)$$

where we replaced the θ -dependent term $L^r(\Omega, \theta; \varepsilon)$ in the first line of (44) with $L^n(\Omega, \bar{\theta} - \varepsilon)$ for all θ in (45). Take an arbitrary $\theta \in [\bar{\theta} - \varepsilon, \bar{\theta}]$. Note that $L^r(\Omega, \theta; \varepsilon) \leq L^n(\Omega, \bar{\theta} - \varepsilon)$ by design of the rule. Then there are two cases to consider.

Case 1: If $L^r(\Omega, \theta; \varepsilon) = L^n(\Omega, \bar{\theta} - \varepsilon)$, then the pointwise variant of the lower bound in (45) is trivially satisfied with equality at any point that falls under Case 1.

Case 2: If $L^r(\Omega, \theta; \varepsilon) < L^n(\Omega, \bar{\theta} - \varepsilon)$, then for this to be an MPE, the government without commitment must weakly prefer choosing $L^r(\Omega, \theta; \varepsilon)$ over $L^n(\Omega, \bar{\theta} - \varepsilon) > L^r(\Omega, \theta; \varepsilon)$:

$$\begin{aligned} & u(c^*(x^r, L^r(\Omega, \theta; \varepsilon), \Omega, \theta), \Omega, \theta) + \beta \mathbb{E}_{\theta'}[V^n(\Gamma(L^r(\Omega, \theta; \varepsilon), \Omega, \theta), \theta')] \\ & \geq u(c^*(x^r, L^n(\Omega, \bar{\theta} - \varepsilon), \Omega, \theta), \Omega, \theta) + \beta \mathbb{E}_{\theta'}[V^n(\Gamma(L^n(\Omega, \bar{\theta} - \varepsilon), \Omega, \theta), \theta')]. \end{aligned} \quad (46)$$

Furthermore, since by Assumption 1 x and ℓ are q-complements in production, we know that $L^r(\Omega, \theta; \varepsilon) < L^n(\Omega, \bar{\theta} - \varepsilon)$ implies that $x^r > x^n$ and thus

$$\begin{aligned} & u(c^*(x^r, L^n(\Omega, \bar{\theta} - \varepsilon), \Omega, \theta), \Omega, \theta) + \beta \mathbb{E}_{\theta'}[V^n(\Gamma(L^n(\Omega, \bar{\theta} - \varepsilon), \Omega, \theta), \theta')] \\ & > u(c^*(x^n, L^n(\Omega, \bar{\theta} - \varepsilon), \Omega, \theta), \Omega, \theta) + \beta \mathbb{E}_{\theta'}[V^n(\Gamma(L^n(\Omega, \bar{\theta} - \varepsilon), \Omega, \theta), \theta')]. \end{aligned} \quad (47)$$

Combining equations (46) and (47), we see that

$$\begin{aligned} & u(c^*(x^r, L^r(\Omega, \theta; \varepsilon), \Omega, \theta), \Omega, \theta) + \beta \mathbb{E}_{\theta'}[V^n(\Gamma(L^r(\Omega, \theta; \varepsilon), \Omega, \theta), \theta')] \\ & > u(c^*(x^n, L^n(\Omega, \bar{\theta} - \varepsilon), \Omega, \theta), \Omega, \theta) + \beta \mathbb{E}_{\theta'}[V^n(\Gamma(L^n(\Omega, \bar{\theta} - \varepsilon), \Omega, \theta), \theta')]. \end{aligned} \quad (48)$$

From the inequality in (48), it follows that the pointwise variant of the lower bound in (45) is satisfied with strict inequality at any point that falls under Case 2.

Combining Cases 1 and 2, we conclude that (45) indeed represents a lower bound on (44). All that remains to be shown is that the value of (45) is strictly positive for small enough $\varepsilon > 0$. To see

that this is the case under the stated assumption of interior lockdown $L^n(\Omega, \theta) \in (0, 1)$, recall that the optimal lockdown is strictly more severe under lack of commitment than under commitment for interior levels of lockdown. This implies that for small enough $\varepsilon > 0$, for all $\theta \in [\bar{\theta} - \varepsilon, \bar{\theta}]$ we have that welfare strictly increases when we replace $L^n(\Omega, \theta)$ by $L^n(\Omega, \bar{\theta} - \varepsilon) < L^n(\Omega, \theta)$, where the strict inequality follows from Assumption 3, which states that $L^n(\cdot)$ is strictly increasing. Since the density $g(\cdot)$ is strictly positive and continuous in a neighborhood below $\bar{\theta}$ by Assumption 3, the interval $[\bar{\theta} - \varepsilon, \bar{\theta}]$ defines a strictly positive probability mass. Combining the last two insights, the expression in (45) is strictly positive for small enough $\varepsilon > 0$.

This concludes the proof that the imposition of such a rule strictly increases welfare. \square

B Details of Quantitative Exercise

B.1 Fundamentals

We study an infinite-horizon economy in discrete time, with periods indexed by $t = 0, 1, 2, \dots$ in the sequence formulation. Next period's value of some current-period variable X is denoted by X' in the recursive formulation.

The government chooses a lockdown policy

$$L \in [0, 1] \tag{49}$$

such that $L = 0$ denotes no lockdown (i.e., everyone goes to work) and $L = 1$ denotes full lockdown (i.e., no one goes to work).

The health state is

$$\Omega = (S, I, R, D) \in [0, 1]^4 \tag{50}$$

such that

$$S + I + R + D = 1. \tag{51}$$

The mass of potential workers, given health state Ω and lockdown policy L , is

$$\tilde{\ell}(\Omega, L) = (1 - L)(S + I + R). \tag{52}$$

The health state dynamics, given health state Ω and lockdown policy L , is

$$\Omega' = \Gamma(\Omega, L). \quad (53)$$

The health state dynamics in recursive formulation, given health state Ω and lockdown policy L , are described by the following system of difference equations:

$$S' = \left[1 - \left(\rho_1 (1 - L)^2 + \rho_2 \right) I \right] S \quad (54)$$

$$I' = \left[1 - \rho_3 - \rho_4 + \left(\rho_1 (1 - L)^2 + \rho_2 \right) S \right] I \quad (55)$$

$$R' = R + \rho_3 I \quad (56)$$

$$D' = D + \rho_4 I. \quad (57)$$

Special attention must be paid to the treatment of corner cases, in which one or more of S' , I' , R' , or D' fall outside of the feasible range $[0, 1]$. In this case, flow rates between all health states (i.e., not just the infeasible health states) need to be adjusted to guarantee $(S', I', R', D') \in [0, 1]^4$.

The health state dynamics in the sequence formulation for $t \geq 1$, given initial health state (S_0, I_0, R_0, D_0) , are given by

$$R_t = R_0 + \rho_3 \sum_{\tau=0}^{t-1} I_\tau \quad (58)$$

$$D_t = D_0 + \rho_4 \sum_{\tau=0}^{t-1} I_\tau \quad (59)$$

If $R_0 = D_0 = 0$, which we assume throughout, then we can combine equations (58) and (59) to get

$$D_t = \frac{\rho_4}{\rho_3} R_t \quad (60)$$

Furthermore, from the adding-up constraint in equation (51) we have

$$S_t = 1 - I_t - R_t - D_t \quad (61)$$

$$= 1 - I_t - \left(1 + \frac{\rho_4}{\rho_3} \right) R_t. \quad (62)$$

Therefore, as long as $R_0 = D_0 = 0$, then we can write the entire problem in terms of the reduced

health state (I_t, R_t) . Note that this formulation implicitly restricts the set of feasible health states (S, I, R, D) .

Factor input prices are given by

$$\text{cost of intermediate inputs (fixed): } r > 0 \quad (63)$$

$$\text{competitive wage (determined in equilibrium): } w > 0. \quad (64)$$

The productivity penalty factor from being infected is

$$\gamma \in [0, 1]. \quad (65)$$

Aggregate economic quantities are as follows:

$$\text{aggregate investment in intermediate inputs: } x \quad (66)$$

$$\text{aggregate effective labor supply: } \ell \leq \bar{\ell}(\Omega, L) \quad (67)$$

$$\text{upper bound on aggregate effective labor supply: } \bar{\ell}(\Omega, L) = (1 - L)(S + \gamma I + R) \quad (68)$$

$$\text{gross output: } y(x, \ell) = Ax^\alpha \ell^{1-\alpha} \quad (69)$$

$$\text{aggregate consumption: } c = w\ell \quad (70)$$

$$\text{aggregate payments to intermediate-input suppliers: } d = rx. \quad (71)$$

Per-capita (alive) economic quantities are as follows:

$$\text{per-capita consumption: } \bar{c} = \frac{c}{S + I + R} \quad (72)$$

$$\text{flow value of being alive: } \nu \in \mathbb{R}. \quad (73)$$

The period utility function is taken to be

$$u(c, \Omega) = (S + I + R)(\ln(\bar{c}) + \nu). \quad (74)$$

The inter-period discount factor is

$$\beta \in [0, 1]. \quad (75)$$

B.2 Problem with Commitment

In the problem with commitment, the period state for all agents is Ω . The firm takes as given lockdown policy each period, which it treats as known. In turn, the government with commitment anticipates that the firm will react to its contemporaneous lockdown policy, which it chooses based on the prevailing health state Ω .

The firm's period profits, given health state Ω and lockdown policy L , are

$$\begin{aligned} \pi(\Omega, L) &= \max_{x, \ell} \left\{ Ax^\alpha \ell^{1-\alpha} - rx - w\ell \right\} & (76) \\ \text{s.t. } & x \geq 0 \\ & \ell \in [0, (1-L)(S + \gamma I + R)] \\ & r, w \text{ given.} \end{aligned}$$

The firm's optimality conditions with respect to investment x and labor ℓ are

$$[\partial x]: \quad r = \alpha Ax^{\alpha-1} \ell^{1-\alpha} \quad (77)$$

$$\implies x = \left(\frac{\alpha A}{r} \right)^{1/(1-\alpha)} \ell \quad (78)$$

$$[\partial \ell]: \quad w = (1-\alpha) Ax^\alpha \ell^{-\alpha} \quad (79)$$

$$\implies w = (1-\alpha) A^{1/(1-\alpha)} \left(\frac{\alpha}{r} \right)^{\alpha/(1-\alpha)}, \quad (80)$$

which shows that wages are invariant to lockdown policy or the health state.

Furthermore, market clearing imposes that

$$\ell = (1-L)(S + \gamma I + R). \quad (81)$$

Aggregate consumption is then given by

$$c = w\ell \quad (82)$$

$$= (1 - \alpha) Ax^\alpha \ell^{1-\alpha} \quad (83)$$

$$= (1 - \alpha) A^{1/(1-\alpha)} \left(\frac{\alpha}{r}\right)^{\alpha/(1-\alpha)} \ell. \quad (84)$$

Per-capita consumption is

$$\bar{c} = \frac{(1 - \alpha) A^{1/(1-\alpha)} \left(\frac{\alpha}{r}\right)^{\alpha/(1-\alpha)} \ell}{S + I + R}. \quad (85)$$

Aggregate payments to intermediate-input suppliers are

$$d = rx \quad (86)$$

$$= \alpha Ax^\alpha \ell^{1-\alpha} \quad (87)$$

$$= \left(\frac{1}{r}\right)^{\alpha/(1-\alpha)} (\alpha A)^{1/(1-\alpha)} \ell. \quad (88)$$

Putting everything together, the government with commitment solves

$$V^e(\Omega) = \max_L \{u(c, \Omega) + \beta V^e(\Omega')\} \quad (89)$$

$$\text{s.t. } L \in [0, 1]$$

$$c = (1 - \alpha) A^{1/(1-\alpha)} \left(\frac{\alpha}{r}\right)^{\alpha/(1-\alpha)} (1 - L) (S + \gamma I + R)$$

$$\Omega' = \Gamma(\Omega, L).$$

B.3 Problem without Commitment

In the problem without commitment, the period state for the firm is Ω , while that for the government is (x, Ω) . The firm anticipates the government's lockdown policy L each period and chooses investment x according to the same no-arbitrage condition as in equation (78):

$$x = \left(\frac{\alpha A}{r}\right)^{1/(1-\alpha)} \ell. \quad (90)$$

Given investment x and lockdown policy L , labor input ℓ is chosen to maximize profits:

$$\begin{aligned} \pi(x, \Omega, L) &= \max_{\ell} \left\{ Ax^{\alpha} \ell^{1-\alpha} - rx - w\ell \right\} \\ \text{s.t. } \ell &\in [0, (1-L)(S + \gamma I + R)]. \end{aligned} \quad (91)$$

This yields the following first-order necessary condition for optimality:

$$[\partial \ell] : \quad w = (1 - \alpha) Ax^{\alpha} \ell^{-\alpha} \quad (92)$$

$$\implies \quad w = (1 - \alpha) A^{1/(1-\alpha)} \left(\frac{\alpha}{r} \right)^{\alpha/(1-\alpha)}, \quad (93)$$

which shows that wages are invariant to lockdown policy or the health state.

Market clearing imposes that

$$\ell = (1 - L)(S + \gamma I + R). \quad (94)$$

However, the government with no commitment treats the firm's investment x as sunk and not affected by its contemporaneous lockdown policy, which it chooses based on the prevailing state (x, Ω) . Mathematically, this means that the no-arbitrage condition in equation (78) still holds but is plugged into the firm optimality condition after taking FOCs, rather than being plugged into the firm's problem before taking FOCs, which would be the case under commitment.

We are looking for a Markov perfect equilibrium in which the firm chooses investment $x^n(\Omega)$ as a function of the prevailing health state Ω and as the best response to the government lockdown policy $L(x^n(\Omega), \Omega)$, which itself is chosen based on the firm's choice of investment $x^n(\Omega)$ and the prevailing health state Ω .

Putting everything together, the government with no commitment solves

$$W^n(x, \Omega) = \max_L \{u(c, \Omega) + \beta V^n(\Omega')\} \quad (95)$$

$$V^n(\Omega') = \max_{L'} \{u(c', \Omega') + \beta V^n(\Gamma(\Omega', L'))\}$$

$$\text{s.t. } L, L' \in [0, 1]$$

$$c = (1 - \alpha) Ax^\alpha [(1 - L)(S + \gamma I + R)]^{1-\alpha}$$

$$c' = (1 - \alpha) A [x^n(\Omega')]^\alpha [(1 - L')(S' + \gamma I' + R')]^{1-\alpha};$$

$x^n(\Omega')$ and $L'(x^n(\Omega'), \Omega')$ form a Markov perfect equilibrium given Ω' :

$$\Omega' = \Gamma(\Omega, L)$$

B.4 Optimal Lockdown Policy with and without Commitment

Then, the FOC for the government with commitment is

$$\frac{d}{dL} [u + \beta V^e] = 0 \quad (96)$$

$$\Leftrightarrow \frac{\partial c}{\partial L} \frac{\partial u}{\partial c} + \beta \frac{dV^e}{dL} = 0 \quad (97)$$

$$\Leftrightarrow \underbrace{-(1 - \alpha) A^{\frac{1}{1-\alpha}} \left(\frac{\alpha}{r}\right)^{\frac{\alpha}{1-\alpha}} (S + \gamma I + R)}_{=\frac{\partial c}{\partial L}} \underbrace{\frac{1}{(1 - \alpha) A^{1/(1-\alpha)} \left(\frac{\alpha}{r}\right)^{\frac{\alpha}{1-\alpha}} (1 - L)(S + \gamma I + R)}}_{=\frac{\partial u}{\partial c}} \quad (98)$$

$$+ \beta \frac{dV^e}{dL} = 0$$

$$\Leftrightarrow \frac{1}{1 - L} = \beta \frac{dV^e}{dL}. \quad (99)$$

In comparison, the FOC for the government with no commitment is

$$\frac{d}{dL} [u + \beta V^n] = 0 \quad (100)$$

$$\Leftrightarrow \frac{\partial c}{\partial L} \frac{\partial u}{\partial c} + \beta \frac{dV^n}{dL} = 0 \quad (101)$$

$$\Leftrightarrow \underbrace{-(1 - \alpha) Ax^\alpha \frac{1 - \alpha}{(1 - L)^\alpha} (S + \gamma I + R)^{1-\alpha}}_{=\frac{\partial c}{\partial L}} \underbrace{\frac{1}{(1 - \alpha) Ax^\alpha [(1 - L)(S + \gamma I + R)]^{1-\alpha}}}_{=\frac{\partial u}{\partial c}} + \beta \frac{dV^n}{dL} = 0 \quad (102)$$

$$\Leftrightarrow \frac{1 - \alpha}{1 - L} = \beta \frac{dV^n}{dL}. \quad (103)$$

From this, we see that the government with no commitment behaves as if it weighs current period utility by a factor $(1 - \alpha) \in (0, 1)$. In other words, the government with no commitment is relatively more patient than the government with commitment.

B.5 Vaccine Arrival and Backward Induction

We assume that in period $T \geq 0$, a vaccine arrives deterministically, preventing any new infections from date T onwards. The same formulation as above applies, but with time entering the problem. Specifically, the health state dynamics are now time-dependent:

$$\rho_{1,t} = \begin{cases} \rho_1 & \text{for } t < T \\ 0 & \text{for } t \geq T. \end{cases} \quad (104)$$

Note that the infections in period T continue to prevail and evolve according to the health state dynamics for $t \geq T$, taking into account $\rho_{1,t}$:

$$\Omega' = \Gamma(\Omega, L, t). \quad (105)$$

Since the lockdown policy L does not affect health state dynamics for $t \geq T$ and stricter lockdowns (i.e., higher values of L) are costly in terms of consumption utility, we know that no lockdown is optimal for $t \geq T$:

$$L^*(\Omega, t) \begin{cases} \in [0, 1] & \text{for } t < T \\ = 0 & \text{for } t \geq T. \end{cases} \quad (106)$$

Following this logic, the dynamic program can be split into two parts. First, consider the problem from date $t \geq T$ onwards, which is after the arrival of the vaccine. Given that no lockdown is optimal for $t \geq T$, the problem of the government with commitment and that with no commitment

coincide and can be written as

$$V^{vacc}(\Omega) = u(c, \Omega) + \beta V^{vacc}(\Omega') \quad (107)$$

$$\text{s.t. } c = (1 - \alpha) A^{1/(1-\alpha)} \left(\frac{\alpha}{r}\right)^{\alpha/(1-\alpha)} (S + \gamma I + R) \quad (108)$$

$$\Omega' = \Gamma(\Omega, 0). \quad (109)$$

Second, consider the problem of the government with or without commitment in period $t < T$, which is before the arrival of the vaccine. Given $V^{vacc}(\Omega)$, we can solve for $V^e(\Omega, t)$, $W^n(x, \Omega, t)$, and $V^n(\Omega, t)$ by backward induction for $t = T - 1, T - 2, \dots, 0$.

B.6 Value of a Statistical Life

We calculate the value of a statistical life, VSL , before the arrival of the pandemic as

$$VSL = \sum_{t=0}^{t^{max}-1} \frac{FVSL}{(1+r)^t} \quad (110)$$

$$= \frac{FVSL \left(1 - \left(\frac{1}{1+r}\right)^{t^{max}}\right)}{1 - \frac{1}{1+r}}, \quad (111)$$

where $t_{max} = 37 \times 52 = 1,924$ is the average number of residual weeks of life and $FVSL$ is the weekly flow value of a statistical life. Therefore, the flow value of a statistical life is

$$FVSL = \frac{VSL \times \left(1 - \frac{1}{1+r}\right)}{1 - \left(\frac{1}{1+r}\right)^{t^{max}}} \quad (112)$$

$$= \frac{VSL \times \frac{r}{1+r}}{1 - \left(\frac{1}{1+r}\right)^{t^{max}}}. \quad (113)$$

To translate the flow value of a statistical life ($FVSL$) into a flow value of being alive (v), we use the standard value of a statistical life calculation (Glover et al., 2020),

$$FVSL = \frac{u(\bar{c}, (1, 0, 0, 0))}{u_c(\bar{c}, (1, 0, 0, 0))} \quad (114)$$

$$= \frac{\ln(\bar{c}) + v}{\bar{c}}, \quad (115)$$

where \bar{c} is the weekly per-capita consumption before the pandemic and v is the flow utility from being alive. Rearranging, we get

$$v = FVSL \times \frac{1}{\bar{c}} - \ln(\bar{c}). \quad (116)$$

Assuming a value of VSL of USD 11.5 million (Greenstone and Nigam, 2020) and a weekly interest rate of $r = (1 + 0.03)^{1/52} - 1$, we have $FVSL = 9,827.09$. Assuming in addition that $\bar{c} = 45,175/52$, as in [Glover et al. \(2020\)](#), we get

$$v = 9,827.09 \times \frac{52}{45,175} - \ln\left(\frac{45,175}{52}\right) \quad (117)$$

$$= 4.54. \quad (118)$$